arXiv:1602.06411v1 [cs.DS] 20 Feb 2016

On the Size and the Approximability of Minimum Temporally Connected Subgraphs

Kyriakos Axiotis and Dimitris Fotakis

School of Electrical and Computer Engineering, National Technical University of Athens, Greece
kaxiotis@corelab.ntua.gr, fotakis@cs.ntua.gr

Abstract

We consider *temporal graphs* with discrete time labels and investigate the size and the approximability of minimum temporally connected spanning subgraphs. We present a family of minimally connected temporal graphs with n vertices and $\Omega(n^2)$ edges, thus resolving an open question of (Kempe, Kleinberg, Kumar, JCSS 64, 2002) about the existence of sparse temporal connectivity certificates. Next, we consider the problem of computing a minimum weight subset of temporal edges that preserve connectivity of a given temporal graph either from a given vertex r (r-MTC problem) or among all vertex pairs (MTC problem). We show that the approximability of r-MTC is closely related to the approximability of Directed Steiner Tree and that r-MTC can be solved in polynomial time if the underlying graph has bounded treewidth. We also show that the best approximation ratio for MTC is at least $O(2^{\log^{1-\varepsilon} n})$ and at most $O(\min\{n^{1+\varepsilon}, (\Delta M)^{2/3+\varepsilon}\})$, for any constant $\varepsilon > 0$, where M is the number of temporal edges and Δ is the maximum degree of the underlying graph. Furthermore, we prove that the unweighted version of MTC is APX-hard and that MTC is efficiently solvable in trees and 2-approximable in cycles.

1998 ACM Subject Classification F.1.3 Complexity Measures and Classes, F.2.2 Nonnumerical Algorithms and Problems, G.2.2 Graph Theory

Keywords and phrases Temporal Graphs, Temporal Connectivity, Approximation Algorithms, Graph Theory

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

Graphs and networks are ubiquitous in Computer Science, as they provide a natural and useful abstraction of many complex systems (e.g., transportation and communication networks) and processes (e.g., information spreading, epidemics, routing), and also of the interaction between individual entities or particles (e.g., social networks, chemical and biological networks). Traditional graph theoretic models assume that the structure of the network and the strength of interaction are time-invariant. However, as observed in e.g., [3, 19], in many applications of graph theoretic models, the availability and the weights of the edges are actually time-dependent. For instance, one may think of information spreading and distributed computation in dynamic networks (see e.g., [6, 12, 19, 20]), of mobile adhoc and sensor networks (see e.g., [22]), of transportation networks and route planning (see e.g., [3, 15]), of epidemics, biological and ecological networks (see e.g., [17, 19]), and of influence systems and coevolutionary opinion formation (see e.g., [5, 9]).

Several variants of time-dependent graphs have been suggested as abstractions of such settings and computational problems (see e.g., [6] and the references therein). No matter the particular variant, the main research questions are usually related either to optimizing or exploiting temporal connectivity or to computing short time-respecting paths (see e.g., [1, 2,

© Kyriakos Axiotis and Dimitris Fotakis;
licensed under Creative Commons License CC-BY
Conference title on which this volume is based on.
Editors: Billy Editor, Bill Editors; pp. 1–26

Leibniz International Proceedings in Informatics
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

3, 13, 15, 19, 21]). In this work, we adopt the simple and natural model of temporal graphs with discrete time labels [19] (and its extension with multiple labels per edge [21]), and study the existence of dense minimally connected temporal graphs and the approximability of temporally connected spanning subgraphs with minimum total weight.

Temporal Graphs and Temporal Connectivity. A *temporal graph* is defined on a time-invariant set of n vertices. Each (undirected) edge e is associated with a set of discrete time labels denoting when e is available. If every edge is associated with a single time label, as in [19], the temporal graph is *simple*. An edge e available at time t comprises a temporal edge (e, t) and there is a positive weight $w(e, t)$ associated with it. A (resp. strict) *temporal* (or *time-respecting*) path is a sequence of temporal edges with non-decreasing (resp. increasing) time labels. So, temporal paths respect the time availability constraints of the edges.

Given a source vertex r, a temporal graph is (temporally) *r-connected* if there is a temporal path from r to any other vertex. A temporal graph is (temporally) *connected* if there exists a temporal path from any vertex to any other vertex. We study the existence of dense minimally connected temporal graphs and the optimization problems of computing a minimum weight subset of temporal edges that preserve either r-connectivity or connectivity. We refer to these optimization problems as (Minimum) Single-Source Temporal Connectivity (or r-MTC, in short) and (Minimum) All-Pairs Temporal Connectivity (or MTC, in short). They arise as natural generalizations of Minimum Spanning Tree and Minimum Arborescence in temporal networks, and to the best of our knowledge, their approximability has not been determined so far (but see [1, 18] for some results on variants or special cases).

Previous Work and Motivation. The model of simple temporal graphs with discrete time labels was introduced in [19]. It is essentially equivalent to the model of scheduled networks [3], where each edge is available in a time interval. [3, 19] investigated how time availability restrictions on the edges affect certain graph properties. Berman [3] presented an algorithm for reachability by temporal paths and proved that an analogue of the max-flow-min-cut theorem holds for temporal graphs. Kempe et al. [19] focused on vertex-disjoint temporal paths and showed that Menger's theorem does not generalize to temporal graphs. They also identified a simple forbidden topological minor for Menger's theorem in temporal graphs. Mertzios at al. [21] introduced multiple labels per edge and studied the number of temporal edges required for a temporal design to guarantee certain graph properties. Interestingly, they proved that a variant of Menger's theorem, which also takes time into account, holds in all temporal graphs. A key technical tool in [3, 19, 21] is the time-expanded version of a temporal graph, which reduces reachability, edge-disjoint path and vertex-disjoint path questions in temporal graphs to similar questions in standard directed graphs.

Our motivation comes from a natural open question in [19, Section 6]. Attempting an analogy between spanning trees of (standard undirected) graphs and connectivity certificates of temporal graphs, Kempe et al. asked whether any simple temporal graph admits a sparse connectivity certificate. They observed that any r-connected temporal graph has a time-respecting arborescence with $n - 1$ edges that serves as a sparse r-connectivity certificate. For general temporal connectivity, however, minimum temporal connectivity certificates of different graphs have different sizes. Kempe et al. observed that an allocation of time labels to the edges of the hypercube makes it minimally temporally connected. Hence, there are temporal graphs on n vertices with temporal connectivity certificates of $\Omega(n \log n)$ edges. An open question in [19, Section 6] was to determine the tightest function $c(n)$ for which any temporally connected graph on n vertices has a temporal connectivity certificate with at most $c(n)$ edges. A trivial upper bound on $c(n)$ is $O(n^2)$, since taking n time-respecting arborescences, each rooted at a different vertex, results in a temporally connected subgraph.

Kempe et al. observed that if we consider strict temporal paths and allow for the same time label at different edges, $c(n) = \Omega(n^2)$ (e.g., consider K_n with the same time label on all edges). Nevertheless, for connectivity with strict temporal paths and distinct time labels, the best known lower bound on $c(n)$ is $\Omega(n \log n)$ ([1], by a labeling of the hypercube).

Contribution. In this work, we resolve the open question of [19] and derive upper and lower bounds on the approximability of Single-Source and All-Pairs Temporal Connectivity.

In Section 3, we construct a family of simple temporal graphs with $3n$ vertices and roughly $n(n+9)/2$ edges which are almost minimally temporally connected, in the sense that the removal of any subset of $5n$ edges results in a disconnected temporal graph[1] (Theorem 3.1). Hence, we show that $c(n) = \Theta(n^2)$ (i.e., there are graphs with dense minimum temporal connectivity certificates), thus resolving the open question of [19]. Our construction is essentially best possible and can be easily extended to connectivity by strict temporal paths (with distinct time labels on the edges). An interesting feature of our construction (and an indication of its tightness) is that slightly increasing the time label of a single temporal edge results in a temporal connectivity certificate with $O(n)$ edges!

Given the huge gap on the size of temporal connectivity certificates, it is natural to ask about the complexity and the approximability of Single-Source and All-Pairs Temporal Connectivity. Previous work shows that we can decide if a temporal graph is connected in polynomial time (see e.g., [1, 3, 19]) and that Single-Source Temporal Connectivity can be solved in polynomial time in the unweighted case. Another interesting observation is that if we use the time-expanded version of a temporal graph for Minimum Temporal Connectivity, the resulting optimization problems are quite similar to Group Steiner Tree problems. In fact, this observation serves as the main intuition behind several of our results.

In Section 4, we show that the polynomial-time approximability of Single-Source Temporal Connectivity (r-MTC) is closely related to the approximability of the classical Directed Steiner Tree problem. Using a transformation from Directed Steiner Tree to r-MTC (Theorem 4.1) and [16, Theorem 1.2], we show that r-MTC cannot be approximated within a ratio of $O(\log^{2-\varepsilon} n)$, for any constant $\varepsilon > 0$, unless NP $\subseteq$ ZTIME$(n^{\text{poly}\log n})$. Our transformation also implies that any $o(n^{\varepsilon})$-approximation for r-MTC would improve the best known approximation ratio of Directed Steiner Tree. On the positive side, using a transformation from r-MTC to Directed Steiner Tree and the algorithm of [7], we obtain a polynomial-time $O(n^{\varepsilon})$-approximation, for any constant $\varepsilon > 0$, and a quasipolynomial-time $O(\log^3 n)$-approximation for r-MTC (Theorem 4.2). We also show that r-MTC can be solved in polynomial time if the underlying graph has bounded treewidth (Theorem 4.3).

In Section 5, we consider the approximability of All-Pairs Temporal Connectivity (MTC). Theorem 4.2 implies an $O(n^{1+\varepsilon})$-approximation for MTC (Corollary 5.1). An approximation-preserving reduction to Directed Steiner Forest and [14, Theorem 1.1] imply a polynomial-time $O((\Delta M)^{2/3+\varepsilon})$-approximation for MTC, where M is the number of temporal edges and Δ is the maximum degree of the underlying graph (Theorem 5.2). If M is quasilinear and Δ is polylogarithmic, we obtain an $O(n^{2/3+\varepsilon})$-approximation. On the negative side, a reduction from Symmetric Label Cover implies that MTC cannot be approximated within a factor of $O(2^{\log^{1-\varepsilon} n})$ unless NP $\subseteq$ DTIME$(n^{\text{poly}\log n})$ (Theorem 5.3, see also [10, Section 4]). We also show that the unweighted version of MTC is APX-hard (Theorem 5.4).

[1] Based on Theorem 3.1, we can easily obtain a family of minimally connected temporal graphs with $\Omega(n^2)$ edges (e.g., we remove temporal edges from the graph, as long as connectivity is preserved). For simplicity and clarity, we avoid presenting a tight (but more complicated) construction of dense minimally connected temporal graphs, and stick to almost minimal graphs in the proof of Theorem 3.1.

In Section 6, we show that MTC can be solved optimally, in polynomial time, if the underlying graph is a tree (Theorem 6.1), and that MTC is 2-approximable if the underlying graph is a cycle (Theorem 6.2, but it is open whether MTC remains NP-hard for cycles).

For clarity, we focus on connectivity by (non-strict) temporal paths. However, all our results can be extended (with small changes in the proofs and with the same approximation guarantees and running times) to the case of connectivity by strict temporal paths.

Comparison to Previous Work. Akrida et al. [1] study connectivity by strict temporal paths. Allocating distinct time labels to the hypercube, they obtain a minimal temporally connected graph with $\Omega(n \log n)$ edges. They also show that any allocation of distinct labels to K_n results in a temporal graph that is not minimally connected. However, they do not give any lower bound on the size of temporal connectivity certificates for K_n. Our Theorem 3.1 improves on the lower bound of [1] from $\Omega(n \log n)$ to $\Omega(n^2)$. [1] also shows that computing the maximum number of edges that are redundant for temporal connectivity is APX-hard.

Huang et al. [18] consider the Single-Source (but not the All-Pairs) version of Minimum Temporal Connectivity in simple scheduled networks [3]. They show that the problem is APX-hard. Using a transformation to Directed Steiner Tree, they show that the approximation guarantees of [7] carry over to Single-Source Temporal Connectivity for scheduled networks. Although the approximation guarantees are the same, the reduction of [18] is slightly different and less general than ours in Theorem 4.2 (which we discovered independently). In addition to the approximability result, we present strong inapproximability bounds for r-MTC and show that it is polynomially solvable for graphs with bounded treewidth.

Erlebach et al. [13] study the problem of computing a shortest exploration schedule of a temporal graph, i.e., a shortest strict temporal walk that visits all vertices. They prove that it is NP-hard to approximate the shortest exploration schedule within a factor of $O(n^{1-\varepsilon})$, for any $\varepsilon > 0$, and construct temporal graphs whose exploration requires $\Theta(n^2)$ steps. Since the notion of exploration schedules is much stronger than (r-)connectivity, their results do not have any immediate implications for r-MTC and MTC (e.g., the $\Theta(n^2)$-explorable graphs of [13, Lemma 4] admit a temporally connected subgraph with $O(n)$ edges).

2 The Model and Preliminaries

Throughout, we let $[k] \equiv \{1, \ldots, k\}$, for any integer $k \geq 1$. An (edge weighted) *temporal graph* $\mathcal{G}(V, E, L)$ with vertex set V, edge set E and lifetime L is a sequence of (undirected edge-weighted) graphs $(G_t(V, E_t, w_t))_{t\in[L]}$, where $E_t \subseteq E$ is the set of edges available at time t and $w_t(e)$ (or $w(e, t)$) is the nonnegative weight of each edge $e \in E_t$. We often write $\mathcal{G}$ or $\mathcal{G}(V, E)$, for brevity. A temporal graph $\mathcal{G}$ is *unweighted* if $w(e, t) = 1$ for all $e \in E_t$ and all $t \in [L]$. For each edge $e \in E_t$, we say that (e, t) is a *temporal edge* of $\mathcal{G}$. For each edge $e \in E$, $L_e = \{t \in [L] : e \in E_t\}$ denotes the set of time units (or *time labels*) when e is available. A temporal graph is *simple* if $|L_e| = 1$ for all edges $e \in E$. We usually let n denote the number of vertices and $M = \sum_e |L_e|$ denote the number of temporal edges of $\mathcal{G}$. For temporal connectivity problems, we can assume wlog. that at least one temporal edge is available in each time unit, and thus, $L \leq M$. The (static) graph $G(V, E)$ is the *underlying graph* of $\mathcal{G}$. We say that $\mathcal{G}$ has some (non-temporal) graph theoretic property (e.g., is a tree, a cycle, a clique, has bounded treewidth) if the underlying graph G has this property.

For a vertex set S, $G[S]$ (resp. $\mathcal{G}[S]$) is the underlying (resp. temporal) graph induced by S. A spanning subgraph $\mathcal{G}'$ of a temporal graph $\mathcal{G} = (G_t(V, E_t, w_t))_{t\in[L]}$ is a sequence of graphs $(G'_t(V, E'_t, w_t))_{t\in[L]}$ such that $E'_t \subseteq E_t$. The total weight of $\mathcal{G}'$ is $\sum_{t\in[L]} \sum_{e\in E'_t} w(e, t)$.

Temporal Connectivity. A *temporal* (or *time-respecting*) path is an alternating sequence

of vertices and temporal edges $(v_1, (e_1, t_1), v_2, (e_2, t_2), \ldots, v_k, (e_k, t_k), v_{k+1})$, such that $e_i = \{v_i, v_{i+1}\} \in E_{t_i}$, for all $i \in [k]$, and $1 \leq t_1 \leq t_2 \leq \cdots \leq t_k$. A temporal path is *strict* if $t_1 < t_2 < \cdots < t_k$. Such a temporal path is from v_1 to v_{k+1} (or a temporal $v_1 - v_{k+1}$ path).

A temporal graph $\mathcal{G}$ is (temporally) *r-connected*, for a given source $r \in V$, if there is a temporal path from r to any vertex $u \in V$. A temporal graph $\mathcal{G}$ is (temporally) *connected*, if there is a temporal path from u to v for any ordered pair $(u, v) \in V \times V$. If all temporal paths are strict, $\mathcal{G}$ is strictly connected (or strictly r-connected). An *(r-)connectivity certificate* of $\mathcal{G}$ is any spanning subgraph of $\mathcal{G}$ that is also (r-)connected.

Single-Source and All-Pairs Temporal Connectivity. Given a temporal graph $\mathcal{G}$ and a source vertex r, the problem of (Minimum) *Single-Source Temporal Connectivity* (r-MTC) is to compute a temporally r-connected spanning subgraph of $\mathcal{G}$ with minimum total weight. The optimal solution to r-MTC is a simple temporal graph whose underlying graph is a tree (see [19, Section 6] and Lemma A.1). Given a temporal graph $\mathcal{G}$, the problem of (Minimum) *All-Pairs Temporal Connectivity* (MTC) is to compute a temporally connected spanning subgraph of $\mathcal{G}$ with minimum total weight.

Approximation Ratio. An algorithm A has *approximation ratio* $\rho \geq 1$ for a minimization problem if for any instance I, the cost of A on I is at most ρ times I's optimal cost.

Directed Steiner Tree and Forest. To understand the approximability of r-MTC and MTC, we use reductions from and to Directed Steiner Tree and the Directed Steiner Forest.

Given a directed edge-weighted graph $G(V, E)$ with n vertices, a source $r \in V$ and a set of k terminals $S \subseteq V$, the Directed Steiner Tree (DST) problem asks for a subgraph of G that includes a directed path from r to any vertex in S and has minimum total weight. The best known algorithm for DST is due to Charikar et al. [7] and achieves an approximation ratio of $O(k^\varepsilon)$, for any constant $\varepsilon > 0$, in polynomial time, and of $O(\log^3 k)$ in quasipolynomial time. On the negative side, [16, Theorem 1.2] shows that DST cannot be approximated within a factor $O(\log^{2-\varepsilon} n)$, for any constant $\varepsilon > 0$, unless NP $\subseteq$ ZTIME$(n^{\text{poly}\log n})$.

Given a directed edge-weighted graph $G(V, E)$ with n vertices and m edges, and a collection $D \subseteq V \times V$ of k ordered vertex pairs, the Directed Steiner Forest (DSF) problem asks for a subgraph of G that contains an $s - t$ path for each $(s, t) \in D$ and has minimum total weight. [14] presents a polynomial-time $O(n^\varepsilon \min\{n^{4/5}, m^{2/3}\})$-approximation for DSF, for any constant $\varepsilon > 0$.

3 A Lower Bound on the Size of Temporal Connectivity Certificates

In this section, we construct an infinite family of simple temporal graphs with $\Theta(n)$ vertices and lifetime $\Theta(n)$ such that any temporal connectivity certificate has $\Omega(n^2)$ edges. Our construction is essentially best possible, since any temporal graph with n vertices and lifetime L admits a connectivity certificate with $O(\min\{n^2, nL\})$ edges.

▶ **Theorem 3.1.** *For any even $n \geq 2$, there is a simple connected temporal graph with $3n$ vertices, $n(n+9)/2 - 3$ edges and lifetime at most $7n/2$, so that the removal of any subset of $5n$ edges results in a disconnected temporal graph.*

Proof sketch. For any even n, we construct a simple connected temporal graph $\mathcal{G}$ with $\Theta(n)$ vertices and $\Theta(n^2)$ edges so that virtually any edge is essential for temporal connectivity.

The Construction. For any even n, $\mathcal{G}$ consists of $3n$ vertices partitioned into three sets $A = \{a_1, \ldots, a_n\}$, $H = \{h_1, \ldots, h_n\}$ and $M = \{m_1, \ldots, m_n\}$, with n vertices each.

The underlying graph $G[A]$ is the complete graph K_n and comprises the *dense* part of the construction with $\Theta(n^2)$ edges. The edges of $G[A]$ are partitioned into $n/2$ edge-disjoint

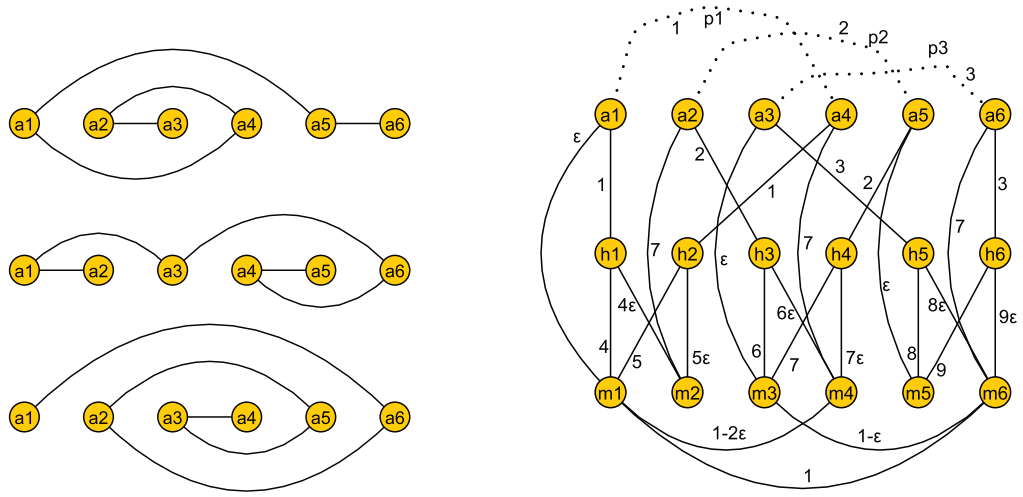

(a) Partition into 3 Hamiltonian paths. (b) Putting the 3 parts together.

Figure 1 The temporal graph constructed in the proof of Theorem 3.1 for $n = 6$.

paths $p_1, \ldots, p_{n/2}$. Each path p_i has length $n-1$ and spans all vertices in A (see Figure 1.a and Lemma B.1). All edges of each path p_i have time label i.

The vertices of H comprise the *intermediate* part of the construction. There are no edges with both endpoints in H. For every $i \in [n/2]$, one endpoint of the path p_i is connected to h_{2i-1} and the other endpoint is connected to h_{2i}. Both edges have time label i.

The vertices of M form the *interconnecting* part of the construction. For each $i \in [n/2]$, we refer to m_{2i-1} (resp. m_{2i}) as the *entry vertex* (resp. the *exit vertex*) for the vertices h_{2i-1} and h_{2i}. There are two edges connecting m_{2i-1} to h_{2i-1} and h_{2i} with labels $n/2 + 2i - 1$ and $n/2 + 2i$, respectively, and two edges connecting m_{2i} to h_{2i-1} and h_{2i} with labels $(n/2 + 2i - 1)\epsilon$ and $(n/2 + 2i)\epsilon$, respectively, for some fixed $\epsilon \in (0, 1/(4n))$. We also connect the vertices of M to each other. For every $i \in [n/2 - 2]$, there are edges connecting m_{2i-1} to m_{2i+2} and to m_n, and a single edge connecting m_{n-3} to m_n. To allocate time labels to these edges, we order them in decreasing order of their endpoint with higher index, breaking ties by ordering them in increasing order of their endpoint with lower index, i.e., the order is $\{m_1, m_n\}$, $\{m_3, m_n\}, \ldots, \{m_{n-3}, m_n\}$, $\{m_{n-5}, m_{n-2}\}$, $\{m_{n-7}, m_{n-4}\}, \ldots, \{m_1, m_4\}$. The time label of the k-th edge in this order is $1 - (k-1)\epsilon$. Finally, for every $i \in [n/2]$, there are an edge with time label ϵ connecting the vertex m_{2i-1} to the vertex a_{2i-1} in A and an edge with time label $n+1$ connecting the vertex m_{2i} to the vertex a_{2i} in A (see also Figure 1.b).

The total number of edges is $n(n+9)/2 - 3$, the number of different time labels is at most $7n/2$, and each edge has a single label (see also Section B.2).

Intuition and Main Claims. The construction is based on the collection $p_1, \ldots, p_{n/2}$ of $n/2$ edge-disjoint paths, where all edges in each path p_i have label i. Extending each path p_i to vertices h_{2i-1} and h_{2i}, we get a path that connects h_{2i} to h_{2i-1} (and vice versa) and to all vertices in A at time i. Moreover, different time labels make these paths essentially independent of each other, in the sense that if a temporal walk begins and ends at time i, it can use only edges with label i (i.e., only edges of this path) to connect h_{2i} to h_{2i-1}. In Section B.4, we formalize this intuition and show that the unique temporal path from h_{2i}

to h_{2i-1} is through path p_i. Therefore, all edges of $G[A]$ must be present in any temporally connected spanning subgraph of $\mathcal{G}$. To achieve a dense underlying graph $G[A]$, we observe that the collection of $n/2$ edge-disjoint paths can be defined so that they go through the same n vertices, in a different order each (see Figure 1.a and Lemma B.1). This describes the main intuition behind our construction and explains how the dense and the intermediate parts work. The only problem now is that H-vertices with high indices, e.g., h_n, cannot reach H-vertices with low indices, e.g., h_1. The vertices in the interconnecting part M serve to carefully connect each h_j to each h_i, with $j > i+1$, without destroying the property that the only temporal path from h_{2i} to h_{2i-1} is through path p_i.

For every vertex pair $h_{2i-1}, h_{2i} \in H$, we introduce a vertex pair $m_{2i-1}, m_{2i} \in M$. As an entry vertex, m_{2i-1} is connected to h_{2i-1} and h_{2i} with "large" labels (larger than $n/2$). Hence, starting from the rest of $\mathcal{G}$, we can reach h_{2i-1} and h_{2i} through m_{2i-1}, but we cannot continue to the edges of p_i (with label $i \leq n/2$). As an exit vertex, m_{2i} is connected to h_{2i-1} and h_{2i} with "very small" labels (at most $1/4$). Thus, starting from m_{2i}, we can reach first h_{2i-1} and h_{2i}, and then all vertices in A and any vertex h_j with index $j > 2i$. Moreover, to avoid creating a temporal path from h_{2i} to h_{2i-1}, the label of the edge $\{h_{2i}, m_{2i-1}\}$ (resp. $\{h_{2i}, m_{2i}\}$) is larger than the label of the edge $\{h_{2i-1}, m_{2i-1}\}$ (resp. $\{h_{2i-1}, m_{2i}\}$).

It remains now to connect the M-vertices to each other, without creating any alternative temporal paths from h_{2i} to h_{2i-1}, for any $i \in [n/2]$. For each $i \in [n/2]$, the edges between M-vertices should create temporal paths from m_{2i-1} and m_{2i} to any vertex m_j with index $j < 2i-1$. On the other hand, they should not create any temporal $m_{2i} - m_{2i-1}$ paths, since then we would have a new temporal $h_{2i} - h_{2i-1}$ path. We introduce roughly n edges between M-vertices and carefully select their "small" labels in $[3/4, 1]$. Furthermore, to achieve temporal connectivity between all vertex pairs, we introduce an edge $\{m_{2i-1}, a_{2i-1}\}$ with the minimum time label ϵ and an edge $\{m_{2i}, a_{2i}\}$ with label $n+1$, for each $i \in [n/2]$.

To complete the proof, in Section B.3, we consider all possible types of ordered vertex pairs and show that the temporal graph $\mathcal{G}$ is indeed connected. Moreover, any subset of at least $5n$ edges includes some edges of $G[A]$. In Section B.4, we show that the removal of any edge from $G[A]$ with label i destroys the unique temporal path from h_{2i} to h_{2i-1}. ◀

We should highlight that increasing the label of edge $\{a_1, m_1\}$ from ϵ to 1, in the graph of Theorem 3.1, results in a temporal graph that admits a connectivity certificate of size $\Theta(n)$ (see Lemma B.2). Moreover, it is not hard to modify the construction of Theorem 3.1 so that all time labels of the edges are distinct, the temporal graph $\mathcal{G}$ is connected by strict temporal paths, and the removal of any subset of $5n$ edges results in a disconnected temporal graph. Therefore, the quadratic lower bound of Theorem 3.1 also applies to connectivity by strict temporal paths and improves on the lower bound of $\Omega(n \log n)$ in [1, Theorem 3].

4 The Approximability of Single-Source Temporal Connectivity

We proceed to study the approximability of the single-source version of Minimum Temporal Connectivity. We show that the polynomial-time approximability of r-MTC is closely related to the approximability of the classical Directed Steiner Tree (DST) problem and that r-MTC can be solved in polynomial-time for graphs of bounded treewidth.

4.1 A Lower Bound on the Approximability of r-MTC

We start with an approximation-preserving transformation from DST to r-MTC. The intuition is that we can use strict temporal paths to "simulate" the directed edges of DST.

▶ **Theorem 4.1.** *Any polynomial-time $\rho(n)$-approximation algorithm for r-*MTC *on simple temporal graphs implies a polynomial-time $\rho(n^2)$-approximation algorithm for* DST.

Proof sketch. We present an approximation-preserving transformation from the DST to r-MTC. Given an instance $I = (G(V, E, w), S, r)$ of DST with $|V| = n$, we construct a temporal graph $\mathcal{G}'$ with n^2 vertices so that (i) any Steiner tree connecting r to S in G can be mapped to an r-connected subgraph of $\mathcal{G}'$ with no larger cost; and (ii) given any r-connected subgraph of $\mathcal{G}'$, we can efficiently compute a feasible Steiner tree for I with no larger cost.

Each vertex $u \in V$ corresponds to a vertex u in the temporal network $\mathcal{G}'$ of I'. For every directed edge $e = (u, v)$ of G, we create $n - 1$ strict temporal $u - v$ paths of length 2. Specifically, for every vertex $u \in V$, $\mathcal{G}'$ contains auxiliary vertices z_i^u, for all $i \in [n-1]$, and temporal edges $\{u, z_i^u\}$ with time label i and weight 0. For every edge $e = (u, v) \in E$, $\mathcal{G}'$ contains temporal edges $\{z_i^u, v\}$ with time label $i+1$ and weight $w(e)$, for all $i \in [n-1]$. Let $Z = \{z_i^u\}_{u \in V, i \in [n-1]}$ be the set of all auxiliary vertices. For every vertex $x \in Z \cup (V \setminus S)$, $x \neq r$, $\mathcal{G}'$ contains a temporal edge $\{r, x\}$ with time label $n + 1$ and weight 0. These edges ensure that r is connected to all non-terminal and auxiliary vertices at no additional cost.

The temporal graph $\mathcal{G}'$ has $\Theta(n^2)$ vertices. Claims (i) and (ii) follow from the construction of $\mathcal{G}'$ and are formally proven in Section C.1. ◀

Directed Steiner Tree cannot be approximated within a ratio of $O(\log^{2-\varepsilon} n)$, for any constant $\varepsilon > 0$, unless NP $\subseteq$ ZTIME$(n^{\text{poly}\log n})$ [16, Theorem 1.2]. Theorem 4.1 implies that this inapproximability result carries over to r-MTC. Moreover, any polynomial-time $o(n^{\varepsilon})$-approximation algorithm for r-MTC would immediately improve the best known approximation ratio of the notoriously difficult DST problem.

4.2 An Approximation Algorithm for r-MTC

In Section C.2, we present an approximation-preserving reduction from r-MTC to DST, thus proving Theorem 4.2 (see also the more general proof of Theorem 5.2). Then, we can use the algorithm of [7] and approximate r-MTC within a ratio of $O(n^{\varepsilon})$, for any constant $\varepsilon > 0$, in polynomial time, and within a ratio of $O(\log^3 n)$ in quasipolynomial time. The reduction of Theorem 4.2 can be easily extended to r-connectivity by strict temporal paths.

▶ **Theorem 4.2.** *Any polynomial-time $\rho(k)$-approximation algorithm for* DST *implies a polynomial-time $\rho(n)$-approximation algorithm for r-*MTC *on general temporal graphs*

4.3 A Polynomial-Time Algorithm for Graphs with Bounded Treewidth

In Section C.3, we prove that r-MTC can be solved in polynomial time, by dynamic programming, if the underlying graph has bounded treewidth (see e.g., [11] about nice tree decompositions and dynamic programming algorithms for graphs of bounded treewidth).

▶ **Theorem 4.3.** *Let $\mathcal{G}$ be a temporal graph on n vertices with lifetime L, source vertex r and treewidth at most k. Then, there is a dynamic programming algorithm which given a nice tree decomposition of G, computes an optimal solution to r-*MTC *in time $O(nk^2 3^k (L + k)^{k+1})$.*

5 The Approximability of All-Pairs Minimum Temporal Connectivity

In this section, we study the approximability of the all-pairs version of Minimum Temporal Connectivity in general temporal graphs. Reducing MTC to r-MTC and to Directed Steiner Forest, we obtain polynomial-time approximation algorithms for MTC, albeit with

not so strong guarantees (Corollary 5.1 and Theorem 5.2). To justify the poor approximation ratios, we reduce Symmetric Label Cover (SLC) to MTC and show that any $\rho(n)$-approximation for MTC implies a $\rho(n^2)$-approximation for SLC (Theorem 5.3). Moreover, using an approximation-preserving reduction from the Steiner Tree problem, we show that the unweighted version of MTC is APX-hard (Theorem 5.4).

5.1 Approximation Algorithms for MTC

Using every vertex of the temporal graph as a source vertex and taking the union of the solutions obtained by the algorithm of Theorem 4.2 for r-MTC, we obtain the following.

▶ **Corollary 5.1.** *For any constant $\varepsilon > 0$, there is a polynomial-time $O(n^{1+\varepsilon})$-approximation algorithm for* MTC *on temporal graphs with n vertices.*

Next, we present a reduction from MTC to Directed Steiner Forest (DSF) that leads to a different algorithm. Although the approximation ratio may be worse than $O(n^{1+\varepsilon})$ in general, this algorithm gives significantly better guarantees if the total number of temporal edges is quasilinear (and if the maximum degree of the underlying graph is polylogarithmic).

▶ **Theorem 5.2.** *Let $\mathcal{G}$ be a temporal graph with n vertices and M temporal edges such that the underlying graph has maximum degree Δ. Then, for any constant $\varepsilon > 0$, there is a polynomial-time $O(M^{\varepsilon} \min\{M^{4/5}, (\Delta M)^{2/3}\})$-approximation algorithm for* MTC *on $\mathcal{G}$. If $M = O(n \operatorname{poly}\log n)$, we obtain an approximation ratio of $O(n^{4/5+\varepsilon})$. If both $M = O(n \operatorname{poly}\log n)$ and $\Delta = O(\operatorname{poly}\log n)$, we obtain an approximation ratio of $O(n^{2/3+\varepsilon})$.*

Proof. The reduction from DSF to MTC is a generalized and refined version of the reduction in Theorem 4.2. Let I be an instance of MTC consisting of an underlying graph $G(V, E)$, a finite set of time labels L_e for each edge e, and a weight $w(e, t)$ for any temporal edge (e, t). We show how to transform I into an instance I' of DSF so that (i) any feasible solution of I can be mapped to a feasible solution of I' with no larger total weight; and (ii) given a feasible solution of I', we can compute a feasible solution of I with no larger total weight.

For convenience, we denote H the edge-weighted directed graph of the DSF instance I'. For every temporal edge (e, t) of $\mathcal{G}$, H contains two vertices $h^1_{(e,t)}$ and $h^2_{(e,t)}$. Intuitively, $h^1_{(e,t)}$ indicates that we may use (e, t) and $h^2_{(e,t)}$ indicates that we actually use (e, t). For each edge $e \in E$, let $l_1(e) < l_2(e) < ... < l_k(e)$ be the time labels in L_e. For every $i \in [k-1]$, H contains a directed edge $(h^1_{(e,l_i(e))}, h^1_{(e,l_{i+1}(e))})$ with weight 0. Intuitively, these edges indicate that we can wait and use e at some later time up to $l_k(e)$. Moreover, for every $i \in [k]$, H contains a directed edge $(h^1_{(e,l_i(e))}, h^2_{(e,l_i(e))})$ with weight $w(e, l_i(e))$. This edge indicates that we actually use the temporal edge $(e, l_i(e))$ and incur the corresponding cost.

For every ordered pair of temporal edges (e_1, t_1), (e_2, t_2) of $\mathcal{G}$, such that $e_1 \neq e_2$, t_2 is the smallest time label in L_{e_2} such that $t_2 \geq t_1$ ($t_2 > t_1$, for strict connectivity), and e_1 and e_2 share a common endpoint, H contains a directed edge $(h^2_{(e_1,t_1)}, h^1_{(e_2,t_2)})$ with weight 0.

For every vertex $v_i \in V$, $i \in [n]$, H contains a pair of terminal vertices s_i and t_i. For every temporal edge (e, t) incident to v_i, H contains a directed edge $(s_i, h^1_{(e,t)})$ with weight 0 and a directed edge $(h^2_{(e,t)}, t_i)$ with weight 0. The set of connection requirements of the DSF instance I' consists of all pairs (s_i, t_j) for all $i, j \in [n]$ with $i \neq j$.

By construction, any temporal $v_i - v_j$ path, which consists of a temporal edge sequence $((e_1, t_1), (e_2, t_2), \ldots, (e_k, t_k))$, with e_1 starting at v_i and e_k ending at v_j, corresponds to a directed $s_i - t_j$ path in H of the form

$$(s_i, h^1_{(e_1,t_1)}), (h^1_{(e_1,t_1)}, h^2_{(e_1,t_1)}), (h^2_{(e_1,t_1)}, h^1_{(e_2,t''_2)}), (h^1_{(e_2,t''_2)}, h^1_{(e_2,t'_2)}),$$

$(h^1_{(e_2,t'_2)}, h^1_{(e_2,t_2)}), (h^1_{(e_2,t_2)}, h^2_{(e_2,t_2)}), \ldots, (h^1_{(e_k,t_k)}, h^2_{(e_k,t_k)}), (h^2_{(e_k,t_k)}, t_j)$

with the same weight and vice versa. Using this observation, we can now establish claims (i) and (ii). Specifically, to show (i), we construct a feasible solution to I' that includes all directed edges of weight 0 and the directed edges $(h^1_{(e,t)}, h^2_{(e,t)})$ corresponding to the temporal edges (e,t) used in the feasible solution to I. Clearly, the two solutions have the same total weight and any temporal $v_i - v_j$ path in the solution to I corresponds to an $s_i - t_j$ path in the solution to I'. To show (ii), we first observe that any directed path from some s_i to some t_j should include some directed edges of the form $(h^1_{(e,t)}, h^2_{(e,t)})$ with weight $w(e,t)$. So, we construct a feasible solution to I that includes the temporal edges (e,t) corresponding to the positive-weight directed edges $(h^1_{(e,t)}, h^2_{(e,t)})$ included in the feasible solution to I'.

In the resulting DSF instance I', the total number of vertices is $O(n+M) = O(M)$ and the number of connection requirements is $O(n^2)$. If the maximum degree of the underlying graph is Δ, the total number of edges is dominated by the edges of the form $(h^2_{(e_1,t_1)}, h^1_{(e_2,t_2)})$, which are $O(\Delta M)$. Applying the approximation algorithm of [14, Theorem 1.1] to the DSF instance I', we obtain a polynomial-time $O(M^\varepsilon \min\{M^{4/5}, (\Delta M)^{2/3}\})$-approximation algorithm, for any constant $\varepsilon > 0$. In the special case where the number of temporal edges is $M = O(n\text{poly}\log n)$, we obtain an $O(n^{4/5+\varepsilon})$-approximation, for any constant $\varepsilon > 0$. If both $M = O(n\text{poly}\log n)$ and the maximum degree of the underlying graph $\Delta = O(\text{poly}\log n)$, we obtain a polynomial-time $O(n^{2/3+\varepsilon})$-approximation algorithm for any constant $\varepsilon > 0$. ◀

5.2 A Lower Bound on the Approximability of MTC

In this section, we present an approximation-preserving reduction from Symmetric Label Cover to MTC. Our reduction along with standard inapproximability results for Symmetric Label Cover indicate that MTC in general temporal networks is hard to approximate.

▶ **Theorem 5.3.** MTC *on temporal graphs with n vertices cannot be approximated within a factor of $O(2^{\log^{1-\varepsilon} n})$, for any constant $\varepsilon > 0$, unless* NP $\subseteq$ DTIME$(n^{\text{poly}\log n})$.

Proof. We present a polynomial-time approximation-preserving reduction from the Symmetric Label Cover (SLC) problem to MTC. In SLC (see e.g., [10, Definition 4.1]), we are given a complete bipartite graph $H(U,W)$, with $|U| = |W|$, a finite set of colors C and a binary relation $R(u,w) \subseteq C \times C$ for every vertex pair $(u,w) \in U \times W$. We seek to assign a color subset $\sigma(u) \subseteq C$ to each vertex $u \in U \cup W$ so that for every vertex pair $(u,w) \in U \times U$, there are colors $a \in \sigma(u)$ and $b \in \sigma(w)$ with $(a,b) \in R(u,w)$ and $\sum_{u \in U \cup W} |\sigma(u)|$, i.e., the total number of colors used, is minimized.

Given an instance of SLC, we create a temporal graph $\mathcal{G}$ whose vertex set V is partitioned into six sets V_U, $V_{C(U)}$, V_W, $V_{C(W)}$, V_X and $\{p,q\}$. There is a correspondence between the vertices of the bipartite graph H and the vertices of $\mathcal{G}$ in the sets V_U and V_W. The vertex sets $V_{C(U)} = V_U \times C$ and $V_{C(W)} = V_W \times C$ serve to encode the color assignment to the vertices of U and W in the SLC instance. Moreover, V_X contains a vertex (u,w,a,b) for every vertex pair $(u,w) \in U \times W$ and every allowable color pair $(a,b) \in R(u,w)$. Intuitively, the vertices of V_X serve to ensure that the color assignment is consistent. Finally, the vertices p and q ensure that the temporal graph $\mathcal{G}$ is connected.

For every $u \in V_U$ and every $(u,a) \in V_{C(U)}$, $\mathcal{G}$ contains a temporal edge $\{u,(u,a)\}$ with time label 1 and weight 1. Similarly, for every $w \in V_W$ and every $(w,b) \in V_{C(W)}$, $\mathcal{G}$ contains a temporal edge $\{w,(w,b)\}$ with time label 4 and weight 1. For every vertex $(u,w,a,b) \in V_X$, $\mathcal{G}$ contains the temporal edges $\{(u,a),(u,w,a,b)\}$ with time label 2 and weight 0 and $\{(u,w,a,b),(w,b)\}$ with label 3 and weight 0. The temporal graph $\mathcal{G}$ contains temporal

edges with label 5 and weight 0 between p and every vertex in $V_U \cup V_{C(U)} \cup V_{C(W)} \cup V_X$ and between q and every vertex in V_U. Moreover, $\mathcal{G}$ contains temporal edges with label 0 and weight 0 between p and every vertex in V_W and between q and every vertex in $V_W \cup V_{C(U)} \cup V_{C(W)} \cup V_X$. We note that the temporal graph $\mathcal{G}$ is connected and has $O(k^2c^2)$ vertices, where $k = |U| = |W|$ and $c = |C|$ (in fact, the number of vertices of $\mathcal{G}$ is of the same order as the total size of all binary relations $R(u, w)$).

We next show that this reduction is approximation-preserving. We first show that any feasible solution to the SLC instance can be mapped to a temporally connected subgraph $\mathcal{G}'$ of $\mathcal{G}$ with at most the same weight. Let us fix any assignment σ of a color set to each vertex of H that is feasible for the SLC instance. We first include in $\mathcal{G}'$ all temporal edges of weight 0. For every vertex $u \in U$ with assigned colors $\sigma(u)$, we include in $\mathcal{G}'$ the temporal edges $\{u, (u, a)\}$, for all $a \in \sigma(u)$. The total weight of these edges is $|\sigma(u)|$. Similarly, for every vertex $w \in W$, we include in $\mathcal{G}'$ the temporal edges $\{w, (w, b)\}$, for all $b \in \sigma(w)$. The total weight of these edges is $|\sigma(w)|$. Therefore, the total weight of the temporal subgraph $\mathcal{G}'$ is equal to the cost of the solution σ for the SLC problem.

It remains to show that $\mathcal{G}'$ is temporally connected. All vertices in $V_U \cup V_{C(U)} \cup V_{C(W)} \cup V_X \cup \{p\}$ are connected with each other (through p) by temporal edges with time label 5. There are also temporal $p-q$ and $q-p$ paths consisting of edges with time label 5 through the vertices of V_U. Similarly, all vertices in $V_W \cup V_{C(U)} \cup V_{C(W)} \cup V_X \cup \{q\}$ are connected with each other (through q) by temporal edges with time label 0. Moreover, there is a temporal path (using edges with time labels 0 and 5) from every vertex in $V_W \cup V_{C(U)} \cup V_{C(W)} \cup V_X \cup \{q\}$ to every vertex in V_U. Also, p is connected to every vertex in V_W with temporal edges of time label 0 and vice versa. All these vertex pairs are connected through temporal paths entirely consisting of 0-weight edges. The really interesting case concerns vertex pairs $(u, w) \in V_U \times V_W$. By the feasibility of the solution σ, for every vertex pair $(u, w) \in U \times W$, there are colors $a \in \sigma(u)$ and $b \in \sigma(w)$ such that $(a, b) \in R(u, w)$. Therefore, the temporal $u-w$ path $(u, (u, a), (u, w, a, b), (w, b), b)$ is included in $\mathcal{G}'$. Hence, $\mathcal{G}'$ is temporally connected.

We also need to show that given a temporally connected subgraph $\mathcal{G}'$ of $\mathcal{G}$, we can efficiently compute an assignment σ of a color set to each vertex in $U \cup W$ that is feasible for the SLC instance and has total cost no larger than the total weight of $\mathcal{G}'$. For every $u \in V_U$ and every temporal edge of the form $(\{u, (u, a)\}, 1)$ included in $\mathcal{G}'$, we include the color a in $\sigma(u)$. Similarly, for every $w \in V_W$ and every temporal edge of the form $(\{w, (w, b)\}, 4)$ included in $\mathcal{G}'$, we include the color b in $\sigma(w)$. Since these are the only edges of $\mathcal{G}$ (and $\mathcal{G}'$) with positive weight, the total cost of σ is equal to the total weight of $\mathcal{G}'$.

It remains to show that σ is a feasible solution to the SLC instance. Let $(u, w) \in U \times W$ of the bipartite graph H in the SLC instance. The crucial observation is that the only way to connect $u \in V_U$ to $w \in V_W$ in $\mathcal{G}'$ is through a temporal path consisting of the temporal edge sequence $(\{u, (u, a)\}, 1)$, $(\{(u, a), (u, w, a, b)\}, 2)$, $(\{(w, b), (u, w, a, b)\}, 3)$, $(\{w, (w, b)\}, 4)$, for some colors a, b such that $(a, b) \in R(u, w)$. This claim immediately implies the feasibility of the assignment σ. To prove this claim, we observe that a temporal $u - w$ path cannot use any temporal edge incident to p or q, since all edges between V_U and $\{p, q\}$ have time label 5 and all edges between V_W and $\{p, q\}$ have time label 0. So, any temporal $u - w$ path in $\mathcal{G}'$ has to move from u to some vertex $(u, a) \in V_{C(U)}$. Such a vertex $(u, a) \in V_{C(U)}$ does not have any neighbors in V_U other than u. Hence, the next vertex of any temporal $u - w$ path in $\mathcal{G}'$ must be to visit some $(u, w, a, b) \in V_X$, where $w \in W$ and $(a, b) \in R(u, w)$. Similarly, since such a vertex $(u, w, a, b) \in V_X$ does not have any neighbors in $V_{C(U)}$ other than (u, a) and any neighbors in $V_{C(W)}$ other than (w, b), we conclude that the next step of any temporal $u - w$ path in $\mathcal{G}'$ must be to the vertex $(w, b) \in V_{C(W)}$. But now, the last

temporal edge used has time label 3, which implies that the $u-w$ path cannot use any edges with time labels 1 and 2 anymore. Thus, the path cannot return to $V_U \cup V_{C(U)}$. The only choice now is that the path moves to w through the temporal edge $(\{w,(w,b)\},4)$, which establishes the claim about the structure of any temporal $u-v$ path in $\mathcal{G}'$.

The discussion above establishes the correctness of the reduction from SLC to MTC. Using the fact that the number of vertices of the temporal graph $\mathcal{G}$ is quadratic in the number of vertices of the bipartite graph H and standard inapproximability results for SLC (e.g., that used in [10]), we conclude the proof of the theorem. ◀

Adjusting the proof of Theorem 5.3, we can get an approximation-preserving reduction from the MinRep problem, which is considered in [8] and is similar to SLC, to MTC. Thus, we obtain that any polynomial-time $\rho(n)$-approximation algorithm for MTC on simple temporal graphs implies a polynomial-time $\rho(n^2)$-approximation algorithm for MinRep. Since the best known approximation ratio for MinRep is $O(n^{1/3}\log^{2/3} n)$ [8, Section 2], any $O(n^{1/6})$-approximation to MTC would imply an improved approximation ratio for MinRep.

5.3 Inapproximability of Unweighted MTC

In Section D, we present an approximation-preserving reduction from the Steiner Tree problem on undirected graphs with edge weights either 1 or 2 to MTC on unweighted temporal graphs, where all temporal edges have weight equal to 1. Since this version of the Steiner Tree problem is known to be APX-hard [4], we obtain the following.

▶ **Theorem 5.4.** MTC *on unweighted temporal graphs is* APX*-hard, and thus it does not admit a PTAS, unless* P = NP*.*

6 All-Pairs Temporal Connectivity on Trees and Cycles

We can do better if the underlying graph is either a tree or a cycle. In Section E.1, Lemma E.1, we show that if the underlying graph is a tree, there is an optimal solution to the MTC problem that uses each edge with at most two time labels. Using this structural property, we can show that MTC can be solved efficiently by dynamic programming if the underlying graph is a tree (see Section E.1 for the details).

▶ **Theorem 6.1.** *Let $\mathcal{G}$ be a temporal tree on n vertices with lifetime L. There is a dynamic programming algorithm that computes an optimal solution to* MTC *on $\mathcal{G}$ in time $O(nL^4)$.*

In Section E.2, we observe that if the underlying graph is a cycle $C_n = (v_0, v_1, \ldots, v_{n-1}, v_0)$, any temporally connected subgraph $\mathcal{G}'$ can be partitioned into *sectors*. A sector is a connected part $(v_i, v_{i+1}, \ldots, v_k)$ of the cycle for which there is a vertex $v_j \notin \{v_i, \ldots, v_{k-1}\}$ such that the temporal paths $p_{\text{incr}} = (v_i, v_{i+1}, \ldots, v_j)$ and $p_{\text{decr}} = (v_k, v_{k-1}, \ldots, v_{j+1})$ are present in $\mathcal{G}'$ (the vertex indices along C_n are taken modulo n). Intuitively, any vertex in the sector $(v_i, v_{i+1}, \ldots, v_k)$ can reach every vertex in C_n through the paths p_{incr} and p_{decr}. Then, in Lemma E.4, we show that there is an optimal solution to the MTC problem on C_n where each edge is shared by at most two different sectors. Then, ignoring edges shared by different sectors and using dynamic programming to determine a near optimal partitioning of C_n into sectors, we obtain the following.

▶ **Theorem 6.2.** *There is a polynomial-time* 2*-approximation algorithm for the* MTC *problem on any temporal cycle C_n.*

References

1. E.C. Akrida, L. Gąsieniec, G.B. Mertzios, and P.G. Spirakis. On Temporally Connected Graphs of Small Cost. In *Proc. of the 13th Workshop on Approximation and Online Algorithms (WAOA '15)*, volume 9499 of *LNCS*, pages 84–96, 2015.
2. E.C. Akrida, L. Gąsieniec, G.B. Mertzios, and P.G. Spirakis. Ephemeral networks with random availability of links: The case of fast networks. *Journal of Parallel and Distributed Computing*, 87:109–120, 2016.
3. K.A. Berman. Vulnerability of scheduled networks and a generalization of Menger's Theorem. *Networks*, 28(3):125–134, 1996.
4. M.W. Bern and P.E. Plassmann. The Steiner Problem with Edge Lengths 1 and 2. *Information Processing Letters*, 32(4):171–176, 1989.
5. K. Bhawalkar, S. Gollapudi, and K. Munagala. Coevolutionary opinion formation games. In *Proc. of the 45th ACM Symposium on Theory of Computing (STOC '13)*, pages 41–50, 2013.
6. A. Casteigts, P. Flocchini, W. Quattrociocchi, and N. Santoro. Time-varying graphs and dynamic networks. *International Journal of Parallel, Emergent and Distributed Systems (IJPEDS)*, 27(5):387–408, 2012.
7. M. Charikar, C. Chekuri, T.-Y. Cheung, Z. Dai, A. Goel, S. Guha, and M. Li. Approximation algorithms for directed steiner problems. *Journal of Algorithms*, 33(1):73–91, 1999.
8. M. Charikar, M. Hajiaghayi, and H.J. Karloff. Improved Approximation Algorithms for Label Cover Problems. *Algorithmica*, 61(1):190–206, 2011.
9. B. Chazelle. The dynamics of influence systems. In *Proc. of the 53rd IEEE Symposium on Foundations of Computer Science (FOCS '12)*, pages 311–320, 2012.
10. Y. Dodis and S. Khanna. Design Networks with Bounded Pairwise Distance. In *Proc. of the 31st ACM Symposium on Theory of Computing (STOC '99)*, pages 750–759, 1999.
11. R.G. Downey and M.R. Fellows. *Fundamentals of Parameterized Complexity*. Texts in Computer Science. Springer, 2013.
12. C. Dutta, G. Pandurangan, R. Rajaraman, Z. Sun, and E.Viola. On the Complexity of Information Spreading in Dynamic Networks. In *Proc. of the 24th ACM-SIAM Symposium on Discrete Algorithms (SODA '13)*, pages 717–736, 2013.
13. T. Erlebach, M. Hoffmann, and F. Kammer. On temporal graph exploration. In *Proc. of the 42nd International Colloquium on Automata, Languages and Programming (ICALP '15)*, volume 9134 of *LNCS*, pages 444–455, 2015.
14. M. Feldman, G. Kortsarz, and Z. Nutov. Improved approximation algorithms for Directed Steiner Forest. *Journal of Computer and System Sciences*, 78(1):279–292, 2012.
15. L. Foschini, J. Hershberger, and S. Suri. On the Complexity of Time-Dependent Shortest Paths. *Algorithmica*, 68(4):1075–1097, 2014.
16. E. Halperin and R. Krauthgamer. Polylogarithmic inapproximability. In *Proc. of the 35th ACM Symposium on Theory of Computing (STOC '03)*, pages 585–594, 2003.
17. P. Holme and J. Saramäki. Temporal Networks. *Physics Reports*, 519(3):97–125, 2012.
18. S. Huang, A.W.-C. Fu, and R. Liu. Minimum Spanning Trees in Temporal Graphs. In *Proc. of the 2015 ACM International Conference on Management of Data (SIGMOD '15)*, pages 419–430, 2015.
19. D. Kempe, J.M. Kleinberg, and A. Kumar. Connectivity and Inference Problems for Temporal Networks. *Journal of Computer and System Sciences*, 64(4):820–842, 2002.
20. F. Kuhn, N.A. Lynch, and R.Oshman. Distributed computation in dynamic networks. In *Proc. of the 42nd ACM Symposium on Theory of Computing (STOC '10)*, pages 513–522, 2010.

21 G.B. Mertzios, O. Michail, I. Chatzigiannakis, and P.G. Spirakis. Temporal network optimization subject to connectivity constraints. In *Proc. of the 40th International Colloquium on Automata, Languages and Programming (ICALP '13)*, volume 7966 of *LNCS*, pages 657–668, 2013.

22 R. O'Dell and R. Wattenhofer. Information dissemination in highly dynamic graphs. In *Proc. of the DIALM-POMC Joint Workshop on Foundations of Mobile Computing (DIALM-POMC '05)*, pages 104–110, 2005.

Appendix

A The Structure of Solutions to Single-Source Temporal Connectivity

The following lemma is implicit in [19, Section 6] and shows that any feasible solution to r-MTC can be transformed to a simple temporal tree without increasing its total weight. Thus, we can always assume that the optimal solution to r-MTC is a simple temporal tree.

▶ **Lemma A.1.** *Given a feasible solution T' for r-MTC, we can obtain a feasible solution T such that (i) T is a simple temporal graph, (ii) the total weight of T does not exceed the total weight of T', and (iii) the set of edges in T form a tree in the underlying graph.*

Proof. We consider any instance of r-MTC and any feasible solution T'. If T' does not include a unique temporal path from r to any other vertex u, we keep the path reaching u earliest and remove all other $r-u$ paths from T'. So, we do not increase the total weight of T' and obtain a feasible solution T that satisfies (ii) and (iii).

As for (i), let $e=\{u,v\}$ be an edge in T. The removal of e partitions the underlying tree into two connected components with vertex sets V_1 and V_2. Let $r,u\in V_1$ and $v\in V_2$. Thus, any temporal path from r to a vertex in V_2 goes through e. Let t_e be the earliest time that a path from r goes from u to v through e. Since there is a temporal $r-v$ path that reaches $v\in V_2$ at time t_e, any path from r to a vertex $v'\in V_2$ can use this path to reach v at t and then proceed to v' using the temporal $v-v'$ path in T. The temporal $v-v'$ path must use time labels at least t_e, since t_e is the earliest time that we reach v from r in T. Thus, we can keep the temporal edge (e,t_e) and remove any other time label associated with e in T. ◀

B Proof of Theorem 3.1: Omitted Technical Details

B.1 Partitioning a Complete Graph into Hamiltonian Paths

We first show that for any even $n\geq 2$, the edges of the complete graph K_n can be partitioned into $n/2$ Hamiltonian paths, each visiting the vertices of K_n in a different order.

▶ **Lemma B.1.** *For any even $n\geq 2$, the edges of the complete graph K_n can be partitioned into $n/2$ Hamiltonian paths.*

Proof. Let $a_0,a_1,\ldots,a_{n-1}$ be the vertices[2] of K_n. For every $i\in\{0,1,...,(n-2)/2\}$, we consider the Hamiltonian path $p_i=(a_i,a_{i-1},a_{i+1},a_{i-2},a_{i+2},\ldots,a_{i-n/2+1},a_{i-n/2-1},a_{i-n/2})$, where all indices are taken modulo n.

[2] For convenience, we index the vertices of K_n from 0 to $n-1$, instead of indexing them from 1 to n, as in the proof of Theorem 3.1. After partitioning the vertices into $n/2$ Hamiltonian paths, we can increase all indices by one and become consistent with the notation in the construction of Theorem 3.1.

We claim that for any $j \neq i$, the paths p_j and p_i do not have any edges in common. To this end, we observe that for any $k \in [n-1]$, the absolute difference between the indices of the k-th vertex and $(k+1)$-th vertex in any such path is exactly k. Therefore, the edge set of each path p_i is uniquely determined by its first vertex a_i. Using this observation inductively, we conclude that for any $j \neq i$, the edge sets of the paths p_j and p_i are disjoint. ◀

B.2 The Number of Edges and the Lifetime of $\mathcal{G}$

The total number of edges in G is $n(n+9)/2 - 3$. Specifically, we have $n(n-1)/2$ edges between vertices in A, $2n$ edges connecting vertices in $H \cup M$ to vertices in A, $2n$ edges connecting vertices in M to vertices in H, and $2(n/2-2)+1$ edges between M-vertices.

The total number of different time labels is at most $7n/2$ and each edge has a single label (so the temporal graph is simple). Specifically, we use $n/2$ labels for the edges in $G[A]$ and for the edges connecting vertices in H to vertices in A, $2n$ labels for the edges connecting vertices in M to vertices in H, 2 labels for the edges connecting vertices in M to vertices in A, and at most $n-2$ labels for the edges in $G[M]$. We also note that the use of non-integral labels in our construction is just for simplicity and without loss of generality.

B.3 The Temporal Graph $\mathcal{G}$ is Connected

We proceed to show that the temporal graph $\mathcal{G}$ constructed in the proof of Theorem 3.1 is indeed connected. To this end, we consider all possible types of ordered vertex pairs (u, v) and show that $\mathcal{G}$ contains a temporal path from u to v. We need to distinguish between several different cases:

$u, v \in A$: We move along the path p_1 of $\mathcal{G}[A]$, where all edges have label 1.

$u \in A$, $v = h_i \in H$: We move from u to a_i using the path p_1, where all edges have label 1. Then, we follow the edge with label $\lfloor i/2 \rfloor$ from a_i to h_i.

$u \in A$, $v = m_{2i-1} \in M$: We first move to h_{2i-1}, as in the previous case, and then follow the edge with label $n/2 + 2i - 1$ from h_{2i-1} to m_{2i-1}.

$u \in A$, $v = m_{2i} \in M$: We move from u to a_{2i} using the path p_1, where all edges have label 1. Then, we follow the edge with label $n+1$ from a_{2i} to m_{2i}.

$u = h_i \in H$, $v \in A$: We move from h_i to a_i using the edge with label $\lceil i/2 \rceil$ and proceed to v along the path $p_{n/2}$, where all edges have label $n/2$.

$u = h_i \in H$, $v = h_j \in H$ and $\lceil i/2 \rceil \leq \lceil j/2 \rceil$: We first move from h_i to a_i, using the edge with label $\lceil i/2 \rceil$, and then on the path $p_{\lceil i/2 \rceil}$ to a_j. Then, we move to h_j using the edge with label $\lceil j/2 \rceil$.

$u = h_i \in H$, $v = h_j \in H$ and $\lceil i/2 \rceil > \lceil j/2 \rceil$: We let $\ell_i = 2\lceil i/2 \rceil$ and $\ell_j = 2\lceil j/2 \rceil$. We first move from h_i to m_{ℓ_i} using the edge with "very small" label $(n/2+i)\epsilon$. We then move to $m_{\ell_i - 3}$, to m_n and to $m_{\ell_j - 1}$ using the edges between M-vertices with "small" labels. Finally, we move from $m_{\ell_j - 1}$ to h_j using the edge with "large" label $n/2 + j$. We note that the labels of the temporal edges between M-vertices are increasing due to the particular ordering that we use.

$u = h_i \in H$, $v = m_{2j-1} \in M$: If $\lceil i/2 \rceil \leq j$, we move from h_i to h_{2j} as in the corresponding case, using edges with labels at most $n/2$. Otherwise, we move from h_i to h_{2j} as in the previous case, using edges with labels at most $n/2 + 2j$. Finally, we move from h_{2j} to m_{2j-1} using the edge with label $n/2 + 2j$.

$u = h_i \in H$, $v = m_{2j} \in M$: We move from h_i to a_{2j} as in the corresponding case, using labels at most $n/2$, and from a_{2j} to m_{2j} using the edge with label $n+1$.

$u = m_{2i-1} \in M$ to any other vertex v: We move from m_{2i-1} to a_{2i-1} using the edge with ϵ and from a_{2i-1} to any other vertex, as in the first four cases, since ϵ is the minimum time label.

$u = m_{2i} \in M$, $v \in A$: We move from m_{2i} to h_{2i}, using the edge with label $(n/2 + 2i)\epsilon$, next to the second endpoint of path p_i, using the edge with label i, and next to any vertex in A, along the path p_i, where all edges have label i.

$u = m_{2i} \in M$, $v = m_{2j} \in M$: We move from m_{2i} to a_{2j} as in the previous case, using labels at most j, and next to m_{2j}, using the edge with label $n + 1$.

$u = m_{2i} \in M$, $v = h_j \in H$: If $i \leq \lceil j/2 \rceil$, we move from m_{2i} to h_{2i} using the edge with label $(n/2 + 2i)\epsilon$, next to the second endpoint of path p_i, using the edge with label i, next to vertex a_j along the path p_i, where where all edges have label i, and finally to h_j using the edge with label $\lceil j/2 \rceil$. Otherwise, we move from m_{2i} to m_{2i-3}, to m_n and to $m_{2\lceil j/2 \rceil - 1}$, using the edges between M-vertices, with "small" labels, and finally from $m_{2\lceil j/2 \rceil - 1}$ to h_j, using the edge with label $n/2 + j$.

$u = m_{2i} \in M$, $v = m_{2j-1} \in M$: If $i \leq j$, we move to h_{2j-1}, as in the first walk of the previous case, using labels at most j, and next to m_{2j-1} using the edge with label $n/2 + 2j - 1$. Otherwise, we follow the second walk of the previous case up to m_{2j-1}. ◀

B.4 Removing a Linear Number of Edges Disconnects $\mathcal{G}$

Since the total number of edges in $\mathcal{G}$ is $n(n + 9)/2 - 3$ and the dense part $\mathcal{G}[A]$ includes $n(n - 1)/2$ of them, removing any subset of at least $5n$ edges results in the removal of at least one edge e between A-vertices. Let us assume that e has label i. We next show that the removal of any edge with label i results in a temporal graph $\mathcal{G}'$ with no temporal path from h_{2i} to h_{2i-1}. Hence, the removal of any set of at least $5n$ edges results in a disconnected temporal graph.

To reach a contradiction, let us assume that there is a temporal path p in $\mathcal{G}'$ from h_{2i} to h_{2i-1}. Since an edge of label i has been removed, the path p_i is not present in $\mathcal{G}'$. Therefore, p should use some edges with labels other than i. We distinguish between two cases:

The first edge of p has label i and the last edge has label larger than i. The only edge incident to h_{2i-1} that can be used as the last edge of p has label $n/2 + 2i - 1$. We note that any edges with labels at most 1 cannot be used by p. The same is true for the edges with label $n + 1$, which connect vertices m_{2j} and vertices a_{2j}, for every $j \in [n/2]$, since all other edges incident to vertices m_{2j} have labels at most 1 and all other edges incident to a_{2j} have labels at most $n/2$. If we ignore all edges with labels at most 1 and $n + 1$, we have no edges between A-vertices and M-vertices and no edges between M-vertices. Therefore, p has to visit vertex a_{2i} first and then to move through the vertices of A. At some point, a vertex $m_j \in M$ must be visited, before p finally reaches h_{2i-1}. But this leads to a contradiction, since there are no usable edges between A-vertices and M-vertices. Hence, no such path p from h_{2i} to h_{2i-1} exists.

Otherwise. The only alternative for p is to begin with label $(n/2 + 2i)\epsilon$ and to end with either label i or label $n/2 + 2i - 1$. Then, the first edge brings p to m_{2i}. At that point, any edge with the minimum label ϵ cannot be used by p. As before, the same is true for the edges with label $n + 1$, which connect vertices m_{2j} and vertices a_{2j}, for every $j \in [n/2]$, since all other edges incident to vertices m_{2j} have labels at most 1 and all other edges incident to a_{2j} have labels at most $n/2$. So, we ignore all edges with labels either ε or $n + 1$. Then, p can only visit vertices m_{2j-1}, with $j < i$, or the vertex m_n (through some edge with a 'small" label in $[3/4, 1]$). From m_n, p cannot proceed any

further, since all edges that connect m_n to some vertex in H have "very small" labels in $(0, 1/4]$. Otherwise, if the last M-vertex visited by p is m_{2j-1}, for some $j < i$, the next edge of p must have either label $n/2 + 2j - 1$ (vertex h_{2j-1}) or label $n/2 + 2j$ (vertex h_{2j}), and then, p cannot proceed any further. Again, we reach a contradiction. Hence, no such path p from h_{2i} to h_{2i-1} exists. ◀

B.5 Tightness of the Construction: Increasing a Single Label

We next show that if in the temporal graph $\mathcal{G}$ constructed in Theorem 3.1, we increase the label of edge $\{a_1, m_1\}$ from ϵ to 1, the resulting temporal graph $\mathcal{G}'$ admits a connectivity certificate with $\Theta(n)$ edges.

▶ **Lemma B.2.** *Let $\mathcal{G}$ be the temporal network constructed in Theorem 3.1 and let $\mathcal{G}'$ be the temporal network obtained from $\mathcal{G}$ if we increase the label of edge $\{a_1, m_1\}$ from ϵ to 1. Then, $\mathcal{G}'$ admits a temporal connectivity certificate with $\Theta(n)$ edges.*

Proof. We show that the temporal graph $\mathcal{G}'$ remains connected if we remove the edges of all paths p_i with $i > 1$. Clearly, this results in a temporal graph with $\Theta(n)$ edges. The proof shows how to modify all temporal $u - v$ paths that are described in Section B.3 and use some edges of a path p_i, with $i > 1$. Again, we need to consider several cases for the initial vertex u and the final vertex v:

- $\boldsymbol{u = h_i \in H, v \in A}$**:** We move from h_i to m_{2i}, to m_{2i-3}, to m_n, and to m_1, using only edges with labels at most 1. Using the edge $\{m_1, a_1\}$, with label 1, we can visit any vertex in A through the path p_1, where all edges have label 1.
- $\boldsymbol{u = h_i \in H, v = h_j \in H}$**:** As in the first case, we move from h_i to the endpoint of the path $p_{\lceil j/2 \rceil}$ which is connected to h_j. This part uses labels at most 1. For the last step, we use the edge from the endpoint of $p_{\lceil j/2 \rceil}$ to h_j, which has label $\lceil j/2 \rceil \geq 1$.
- $\boldsymbol{u = m_{2i} \in M, v \in A}$**:** Similarly to the first case, we move from m_{2i} to m_{2i-3}, to m_n, and to m_1, using only edges with labels at most 1. Using the edge $\{m_1, a_1\}$, with label 1, we can visit any vertex in A through the path p_1, where all edges have label 1.
- $\boldsymbol{u = m_{2i} \in M, v = h_j \in H}$**:** As in the previous case, we from m_{2i} to the endpoint of the path $p_{\lceil j/2 \rceil}$ which is connected to h_j. This part uses labels at most 1. For the last step, we use the edge from the endpoint of $p_{\lceil j/2 \rceil}$ to h_j, which has label $\lceil j/2 \rceil \geq 1$.

◀

C Single-Source Temporal Connectivity

C.1 Proof of Theorem 4.1: Omitted Technical Details

We next prove that (i) any feasible solution T of the DST instance $I = (G(V, E, w), S, r)$ can be mapped to a temporally r-connected subgraph T' of $\mathcal{G}'$ with no larger cost; and (ii) given a temporally r-connected subgraph T' of $\mathcal{G}'$, we can efficiently compute a directed Steiner tree T which connects r to all terminal vertices in S and has no larger cost than T'.

We first show (i). Let T be any feasible solution to DST. For every vertex u, let $d(u)$ be the number of edges on the $r - u$ path in T (we consider the $r - u$ path with the smallest number of edges, if there are many). To transform T into an r-connected subgraph T' of $\mathcal{G}'$, we first include in T' all 0-cost edges incident to r. Next, for each directed edge (u, v) in T, we include in T' the temporal path $u - z^u_{d(u)} - v$, with time labels $d(u)$ and $d(u) + 1$ and weight $w(e)$. Clearly, the total weight of T' does not exceed the total weight of T. Moreover,

there is a temporal path from r to any vertex u of $\mathcal{G}'$. If $u \in Z \cup (V \setminus S)$, we use the direct edge $\{r, u\}$ with time label $n+1$. If $u \in S$, there is a directed $r-u$ path in T. For the i-th edge of this path, $i \leq n-1$, we have added a temporal path of length 2 with labels i and $i+1$. The union of these paths defines a temporal path from r to u.

We next show (ii). Let T' be any r-connected subgraph of $\mathcal{G}'$. Without increasing the cost of T', we can assume that any edge $\{z_i^u, v\}$ with positive weight $w(u,v)$ in T' belongs to a temporal $u-v$ path of length 2, where $u, v \in V$. To transform T' into a feasible solution T of the DST instance I, we include in T the edge $(u,v) \in E$ of any temporal path $u - z_i^u - v$ in T'. Clearly, the total weight of T does not exceed the total weight of T'. To show that T is a feasible solution, we observe that all edges incident to a terminal vertex $u \in S$ have time labels at most n. Therefore, the maximum label in any temporal $r-u$ path is at most n. Hence, for any $u \in S$, T' includes a temporal $r-u$ path consisting of temporal paths of length 2, i.e., of paths of the form $v_k - z_i^{v_k} - v_{k+1}$. Each of these paths becomes a directed edge (v_k, v_{k+1}) in T. These directed edges together form a directed $r-u$ path in T. ◀

C.2 The Proof of Theorem 4.2

Let I be an instance of r-MTC consisting of an underlying graph $G(V,E)$, a source $r \in V$, a finite set of time labels L_e for each edge e, and a cost $w(e,t)$ for any temporal edge (e,t). We show how to transform I into an instance $I' = (H, S, r')$ of DST so that (i) any feasible solution of I can be mapped to a feasible solution of I' with no larger cost; and (ii) given a feasible solution of I', we can compute a feasible solution of I with no larger cost.

The directed graph H of the DST instance contains a non-terminal vertex for each temporal edge (e,t) of $\mathcal{G}$. For every (ordered) pair of temporal edges (e_1,t_1), (e_2,t_2) of $\mathcal{G}$, such that $e_1 \neq e_2$, $t_1 \leq t_2$ (or $t_1 < t_2$, for strict r-connectivity), and e_1 and e_2 share a common endpoint, there is a directed edge $((e_1,t_1),(e_2,t_2))$ with weight $w(e_2,t_2)$ in H. H also contains a root vertex r' corresponding to the source r. For every temporal edge (e,t) incident to r in $\mathcal{G}$, there is a directed edge $(r',(e,t))$ with weight $w(e,t)$ in H. For every vertex $u \in V \setminus \{r\}$, H contains a terminal vertex u'. The terminal set S of I' consists of all these terminal vertices u'. Moreover, for every temporal edge (e,t) incident to u in $\mathcal{G}$, H contains a directed edge $((e,t),u')$ with weight 0. The directed graph H has $O(M)$ vertices, $n-1$ of which are terminal vertices.

We first show (i). Let T be any feasible solution of the r-MTC instance I. By Lemma A.1, we can assume that T is a simple temporal graph and that its underlying graph is a tree (otherwise, we can obtain such a solution from T without increasing its total weight). Thus, T contains a unique temporal path from r to any vertex $u \in V \setminus \{r\}$. To obtain a feasible solution T' for the DST instance I', we do the following, starting from r, in a BFS order: For every pair of vertices $u, v \in V$ such that u is the immediate predecessor of v on the unique temporal $r-v$ path in T, u is reached through the temporal edge (e_1,t_1), and v is reached from u through the temporal edge (e_2,t_2), we add the directed edge $((e_1,t_1),(e_2,t_2))$ to T' (this edge exists in H, because both e_1 and e_2 have u as one of their endpoints and $t_1 \leq t_2$). For every vertex $v \in V$ reachable in T directly from r through the temporal edge (e,t), we add the directed edge $(r',(e,t))$ to T'. Moreover, we add all directed edges of weight 0 to T'. Clearly, the cost of T' does not exceed the cost of T. As for the feasibility of T', we observe that any temporal $r-u$ path in T, which consists of a temporal edge sequence $((e_1,t_1),(e_2,t_2),\ldots,(e_k,t_k))$, with e_1 starting at r and e_k ending at u, corresponds to a directed $r'-u'$ path $(r',(e_1,t_1),(e_2,t_2),\ldots,(e_k,t_k),u')$ in T'

We next show (ii). Let T' be any feasible solution of the DST instance I'. Thus, T' includes a directed path from r' to any vertex $u' \in S$. For each Steiner vertex (e,t) used by

some path in T', we add the corresponding temporal edge to the solution T of the r-MTC instance I. Clearly, the total cost of T is no more than the total weight of T'. We also need to show that T is indeed a feasible solution for the r-MTC instance I. For every vertex $u \in V \setminus \{r\}$, there is a directed path $(r', (e_1, t_1), (e_2, t_2), \ldots, (e_k, t_k), u')$ in T'. By the construction of I', $t_1 \leq t_2 \leq \cdots \leq t_k$ and $(e_1, e_2, \ldots, e_k)$ form an $r - u$ path in the underlying graph of $\mathcal{G}$. Therefore, since the temporal edges $(e_1, t_1), (e_1, t_2), \ldots, (e_k, t_k)$ are present in T, T includes a temporal $r - u$ path for every vertex $u \in V \setminus \{r\}$. ◀

C.3 The Proof of Theorem 4.3

In this section, we present an exact polynomial-time algorithm for r-MTC on temporal graphs with treewidth at most k. The algorithm is based on dynamic programming and its time complexity is $O(nk^2 3^k (L + k)^{k+1})$. To this end, we consider a temporal graph $\mathcal{G}(V, E)$ on n vertices with lifetime L, source vertex let r and treewidth k. Below, we use the standard notation and terminology employed in dynamic programming algorithms for graph theoretic problems on graphs of bounded treewidth (see e.g., [11, Chapter 10]).

We assume a *nice tree decomposition* of $\mathcal{G}$ with $O(nk)$ bags, where the root bag consists only of the vertex r. Also, for each bag i, consider an enumeration $v_1^i, v_2^i, \ldots, v_p^i$ of the vertices in bag i, where $p \leq k + 1$. For each bag i, we define $f(i, a_1, t_1, \ldots, a_{k+1}, t_{k+1})$, where $a_j \in \{0, 1\}$, $t_j \in [L + k + 1]$, as follows:

- Let H be the subgraph of $\mathcal{G}$ induced on the vertices of bag i and all of its descendants in the tree decomposition.
- Let T be a minimum cost temporal subgraph of H such that for every vertex v of H, there exists a vertex v_j^i with $a_j = 1$ and a temporal path from v_j^i to v in T, starting no sooner than t_j.
- Moreover, if $v = v_l^i$ for some l, then this temporal path ends at time t_l.

Note that H is always a temporal forest. Intuitively, $a_j = 1$ means that the j-th vertex of the bag is currently connected to the root, while $a_j = 0$ means that it is not connected yet, and we should connect it via some temporal path in H. Note that some of the arguments of f can be undefined, as some bags may have less than $k + 1$ vertices. The cost of the optimal solution to r-MTC is then $f(1, 1, 1)$, since we start at bag 1 (which only contains the source r) and we are at r at time 1. To break the symmetry introduced by same-label components, for each set of vertices in bag i with the same time label t, we introduce an *ordering* between these vertices. This is done to avoid cyclic dependencies between parent links, which would lead to an infeasible solution. To express such an ordering $x_1, \ldots, x_k$, suppose that the vertices $v_{x_1}^i, \ldots, v_{x_k}^i$ all share the same time label t (the temporal path from r to each of them ends at time t). Then, we set $t_{x_1} = t, t_{x_2} = L + x_1, \ldots, t_{x_k} = L + x_{k-1}$. Thus, we express the required information with space $(L + k)^{k+1}$.

Below, we let $c_1(i)$ and $c_2(i)$ denote the left and right children of bag i respectively. The recurrence relation for computing f is the following:

Case 1: i is an introduce vertex. Suppose wlog that v_p^i is introduced. If $a_p = 0$, then there must exist a temporal path from some vertex v_j^i with $a_j = 1$ to v_p^i in T, the second-to-last vertex of the path being in bag i. This is because v_p^i exists in bag i but none of its descendants, so this is the last chance to connect it to vertices with $a_j = 1$. It is also the last chance to connect vertices with $a_j = 0$ directly to v_p^i (as its children). To this end, we try all possible vertices as parents of v_p^i, and also all possible subsets of vertices with $a_j = 0$ as the set of children of v_p^i. Of course, we adhere to the time labels of the

vertices, as well as the ordering introduced above. Thus, we have that

$$f(i,a_1,t_1,\ldots,0,t_p)=\min_{j\neq i}\Big\{f(c_1(i),a_1,t_1,\ldots,1,t_p)+w_{\text{time}(p)}(v_j^i,v_p^i)\ :\\ \text{time}(j)<\text{time}(p)\vee\\ (\text{time}(j)=\text{time}(p)\wedge j \text{ is before } p \text{ in the ordering})\Big\}$$

where time(j) is the time label that the temporal path from r to v_j^i ends. (If $t_j \leq L$, then time$(j) = t_j$, otherwise time(j) = time$(t_j - L)$, where $t_j - L$ is the index of the predecessor of v_j^i in the ordering of vertices with time(k) = time(j)). Similarly,

$$f(i,a_1,t_1,\ldots,1,t_p)=\min_{S}\Bigg\{f(c_1(i),a_1^*,t_1,\ldots,a_{p-1}^*,t_{p-1})+\sum_{v_j^i\in S}w_{\text{time}(j)}(v_p^i,v_j^i)\ :\\ S\subseteq\{v_j^i\in\{v_1^i,\ldots,v_{p-1}^i\}:a_j=0\wedge\\ (\text{time}(p)<\text{time}(j)\vee\\ (\text{time}(p)=\text{time}(j)\wedge p \text{ is before } j \text{ in the ordering}))\}\Bigg\}$$

where $a_j^* = 0$ if $a_j = 0$ and $v_j^i \notin S$, and $a_j^* = 1$ otherwise. Therefore, the complexity of computing all values of f that fall to this case is $O(nk^2 3^k (L+k)^{k+1})$. This breaks down into $O(nk)$ for the size of the tree, $O(k)$ for iterating over vertices of S and computing time(j), $O(3^{k+1})$ for choices over vertices with $a_j = 1$ as well as vertices with $a_j = 0$ and $v_j^i \in S$, and $(L+k)^{k+1}$ for choices of t_j.

Case 2: i is a forget vertex. Suppose wlog that $v_{p+1}^{c_1(i)}$ is forgotten, and iterate over all possible values t_{p+1}. v_{p+1}^i is either assigned a new time label, or is inserted in the ordering between vertices with the same time label. We have that

$$f(i,a_1,t_1,\ldots,a_p,t_p)=\min_{t_{p+1}\in[L+k]}\Big\{f(c_1(i),a_1,t_1,\ldots,a_{j'},L+p+1,\ldots,a_p,t_p,0,t_{p+1})\ :\\ t_{p+1}\leq L\vee t_{p+1}-L=j\in[p]\Big\}$$

where $v_{j'}^i$ was the successor of v_j^i (if it had one) in the ordering, which now becomes the successor of v_{p+1}^i in the case that $t_{p+1} - L = j$ (v_{p+1}^i is inserted between v_j^i and $v_{j'}^i$ in the ordering). So, in these cases, the total time complexity is $O(nk^2 2^k (L+k)^{k+1})$.

Case 3: i is a join vertex. Then some vertices with $a_j = 0$ will be connected in $c_1(i)$ and some others in $c_2(i)$. Suppose that $P = \{v_j^i \in \{v_1^i,\ldots,v_p^i\} : a_j = 0\}$. Therefore, trying every possible partition of the vertices with $a_j = 0$, we have that

$$f(i,a_1,t_1,\ldots,a_p,t_p)=\min_{S_1,S_2}\Big\{f(c_1(i),a_1^*,t_1,\ldots,a_p^*,t_p)+f(c_2(i),a_1^{**},t_1,\ldots,a_p^{**},t_p)\ :\\ S_1\cup S_2=P \text{ and } S_1\cap S_2=\emptyset\Big\}$$

Here we have that $a_j^* = 0$ and $a_j^{**} = 1$, if $a_j = 0$ and $v_j^i \in S_2$, that $a_j^* = 1$ and $a_j^{**} = 0$, if $a_j = 0$ and $v_j^i \in S_1$, and that $a_j^* = a_j^{**} = 1$, if $a_j = 1$.
Here, the total time complexity is $O(nk^2 3^k (L+k)^{k+1})$. We observe that the vertex order is introduced in the recurrence due to Case 3, since without such an ordering, it would not be possible to solve the two subproblems independently.

Overall we can compute f, and the corresponding minimum cost r-connected temporal subgraph, in time $O(nk^2 3^k (L+k)^{k+1})$. ◀

D APX-hardness of Unweighted MTC: The Proof of Theorem 5.4

In this section, we show that MTC on unweighted temporal graphs is APX-hard and does not admit a PTAS, unless P = NP, using an approximation-preserving reduction from Steiner Tree on undirected graphs with edge weights 1 and 2 to MTC.

Given an undirected edge-weighted graph $G(V, E)$ and a set of terminals $S \subseteq V$, the Steiner Tree (ST) problem asks for a connected subgraph of G that spans S and has minimum total weight. It is easy to see that the optimal solution is always a tree. If the weight of each edge is either 1 or 2, we get the Steiner Tree problem with weights 1 and 2 (ST$(1, 2)$).

In [4], ST$(1, 2)$ is shown APX-hard by a reduction from Vertex Cover on graphs with bounded maximum degree. If n and m are the number of vertices and edges of the Vertex Cover instance respectively, then the ST$(1, 2)$ instance has m terminals and n non-terminals. Furthermore, all the terminals have degree 2 and the graph is bipartite, the one part consists of the terminal vertices and the other part consists of the non-terminal vertices. This means that the total number of edges is at most $2m$. In order to use uniform weights for the edges of the temporal graph, we substitute each edge of weight 2 with a 2-edge path, where each edge has weight 1 and the newly added vertex is a non-terminal vertex. Since at most one non-terminal vertex per edge is added, the resulting number of non-terminal vertices is at most $n + 2m \leq 3m$, assuming that the initial graph is connected and not a tree. So we obtain an instance of the Steiner Tree problem with uniform edge-weights, a set S of m terminal vertices and a set T of at most $3m$ non-terminal vertices.

To reduce this special case of the Steiner Tree problem to MTC with uniform temporal edge weights, we create a temporal variant of the above graph, where each existing edge has label 3. To ensure temporal connectivity, we add some new elements to our temporal graph. We add new vertices p and q and connect every non-terminal vertex in T to p with time label 4 and to q with time label 2. These edges ensure that there is a temporal path from every vertex in $S \cup T$ (terminal and non-terminal vertices, respectively) to every other vertex in the same set.

It remains to ensure temporal connectivity between vertices p and q and to the rest of the temporal graph. Suppose we enumerate the elements of S as $u_1, u_2, \ldots, u_m$. We add m new vertices $a_1, a_2, \ldots, a_m$, composing vertex set A. To create a path from p to any vertex of S, add temporal edges with label 1 from p to every vertex in A, and for each $i \in [m]$, add a temporal edge between a_i and u_i with label 5. Similarly, to connect S to q, add a new vertex set B consisting of vertices $b_1, b_2, \ldots, b_m$, for every i add a temporal edge between u_i and b_i with label 1, and add an edge between q and every vertex in B with label 5. Finally, to create a temporal path from p to q, add a new vertex x, which is connected to p with label 1 and to q with label 5. Then, we solve MTC on the resulting temporal graph $\mathcal{G}$.

We are ready to show the following. Note that Theorem 5.4 follows from Lemma D.1 below and the fact that Steiner Tree with edge weights 1 and 2 is APX-hard.

▶ **Lemma D.1.** *An $(1+\varepsilon)$-approximate solution to* MTC *implies an $(1+12\varepsilon)$-approximate solution to the Steiner Tree instance with uniform weights.*

Proof. The proof consists of two main claims. The first claim is that any feasible solution to MTC on the resulting temporal graph $\mathcal{G}$ uses all the edges that are not in the Steiner Tree instance and that none of them can be part of some temporal path between vertices of S. The second claim is that taking all the edges not in the Steiner Tree instance together with the set of temporal edges corresponding to a feasible solution to the Steiner Tree instance results in a temporally connected subgraph of $\mathcal{G}$.

Having established these claims and assuming that the number of the additional (not in the Steiner Tree instance) edges is $k \leq 4m + 6m + 2 \leq 11m \leq 11\,\mathrm{OPT}$, where OPT is the optimal solution to the Steiner Tree instance, we obtain that the MTC solution has cost at most $(1+\varepsilon)(\mathrm{OPT} + k)$. Deleting these k edges, we have a feasible solution to Steiner Tree of cost at most $(1+\varepsilon)\mathrm{OPT} + \varepsilon k \leq (1 + 12\varepsilon)\mathrm{OPT}$.

Let us first show that any feasible solution to MTC on the temporal graph $\mathcal{G}$ resulting by the reduction above uses all the edges that are not in the Steiner Tree instance. We examine several cases:

Edge between u_i and a_i is missing for some i: Then there is no temporal path from u_i to a_i, since a_i is only adjacent to a temporal edge with label 1, and there is no 1-labeled path from u_i to a_i.

Edge between u_i and b_i is missing for some i: Then there is no temporal path from b_i to u_i, since b_i is only adjacent to a temporal edge with label 5, and there is no 5-labeled path from b_i to u_i.

Edge between a_i and p is missing for some i: Then there is no temporal path from a_i to p, since a_i is only adjacent to a temporal edge with label 5, and there is no 5-labeled path from a_i to p.

Edge between b_i and q is missing for some i: Then there is no temporal path from q to b_i, since b_i is only adjacent to a temporal edge with label 1, and there is no 1-labeled path from q to b_i.

Edge between p and $t \in T$ is missing: Then there is no temporal path from p to t: Going from p to A or x doesn't help, since we then have to use label 5, which is a dead-end. Using another edge with label 4 is also a dead-end.

Edge between q and $t \in T$ is missing: Then there is no temporal path from t to q: Reasons similar as above.

Edge between p and x or edge between q and x is missing: Then there is no temporal path from x to p, or from q to x respectively.

So, we have shown that all temporal edges not present in the Steiner Tree instance are necessary for a feasible solution to the MTC instance. Furthermore, every temporal path between two vertices of S can consist only of temporal edges with label 3: If an edge between S and A is traversed, then we can no longer proceed, since it has label 5. If an edge between S and B is traversed, then an edge between B and q must be traversed, but then we can no longer proceed to S, since we have traversed an edge with label 5. This means that the edges with label 3 in some feasible solution define a feasible solution to the Steiner Tree instance.

Let us now show that taking all the edges not in the Steiner Tree instance together with the set of temporal edges corresponding to a feasible solution to the Steiner Tree instance results in a temporally connected subgraph of $\mathcal{G}$. For this, it suffices to prove that the vertices in $S \cup T \cup \{p\} \cup \{q\}$ are connected to each other with temporal paths, since the rest of the vertices are connected to them with both minimum and maximum labels (1 and 5). Hence, temporal connectivity is implied. We already have that there is a 3-labeled path between each pair of vertices in S. Consider the following cases:

From $u_i \in S$ to $t \in T$: Move from u_i to some $t' \in T$ with label 3, and then to t with label 4 via p. (This also connects u_i to p, and t' to t).

From $u_i \in S$ to q: (u_i, b_i, q)

From $t \in T$ to $u_i \in S$: Move from t to some t' (with label 2 via q) that is connected to $u_j \in S$ with label 3, and then move on the 3-labeled path to u_i. (This also connects t to q, q to t and q to u_i).

From $t \in T$ to p: Move on edge with label 4. (This also connects p to t).
From p to $u_i \in S$: (p, a_i, u_i)
From p to q: (p, x, q)
From q to p: (q, t, p), for some $t \in T$.

Since we have shown all the required facts, the claim is proven. ◀

E Temporal Connectivity on Trees and Cycles: Omitted Proofs

E.1 Temporal Connectivity on Trees: A Polynomial-Time Algorithm

MTC can be solved exactly in polynomial time if the underlying graph is a tree.

▶ **Lemma E.1.** *Let $\mathcal{G}$ be a temporal tree. Then, there is an optimal solution to* MTC *on $\mathcal{G}$ that uses at most two time labels of each edge.*

Proof. Let T be the underlying tree of $\mathcal{G}$. For some edge e, denote by T_A, T_B the connected components in which the T is partitioned after the deletion of e. For some fixed node u in T_A and every node v in T_B, there is a temporal path in the optimal solution that connects u to v. Among all these temporal paths, let $(u, (e_1, t_1), (e_2, t_2), \dots, (e, t), \dots, v)$ be the path with the minimum traversal time for edge e. Note that for all nodes v' of T_B other than v, there exists a path $(u, (e'_1, t'_1), (e'_2, t'_2), \dots, (e, t'), p', v')$ from u to v'. Since $t' \geq t$, the path $(u, (e_1, t_1), (e_2, t_2), \dots, (e, t), p', v')$ is also a valid temporal $u - v'$ path. This means that in some optimal solution to MTC, one time label $t(e, u)$ on edge e can be used by all paths from u to vertices in T_B. If we keep as the time label of e the maximum time label $t(e, v)$, over all vertices $v \in T_A$, this can be used by any temporal path from some vertex in T_A to some vertex in T_B. Hence, for all paths from T_A to T_B, one label suffices for e. By symmetry, the same holds for all paths from T_B to T_A. Overall, there is an optimal solution to MTC where at most two time labels are used for any edge e. ◀

Now, MTC on temporal trees can be solved as follows: We root the tree arbitrarily at some vertex r. For some vertex u, let $f(u, t_i, t_o)$ denote the minimum cost of MTC on the temporal tree T_u rooted at u, with the additional constraint that for every node v in T_u, there must exist a temporal path from u to v using only temporal edges with time label at least t_i and a temporal path from v to u using only temporal edges with label at most t_o. Denote by c_j, for $i \in [1, x_u]$ the children of u in the tree, where x_u is the number of children. Also, let $g(u, j, t_i, t_o)$ denote the minimum cost for MTC on the temporal tree rooted at u but with no children other than $c_1, \dots, c_j$, with the constraints concerning t_i, t_o as stated above, but restricted to the subtrees of the first j children. Obviously, $g(u, 0, t_i, t_o) = 0$ and $g(u, x_u, t_i, t_o) = f(u, t_i, t_o)$. We now observe that

$$g(u, j, t_i, t_o) = \min_{t'_i, t'_o \in \mathbb{N}} \Big\{ g(u, j-1, \max(t_i, t'_o), \min(t_o, t'_i)) + f(c_j, t'_i, t'_o) + q(u, c_j, t'_i, t'_o) \ : \\ t'_i \geq t_i \text{ and } t'_o \leq t_o \Big\}$$

In the equation above, t'_i denotes the time that the edge $\{u, c_j\}$ is traversed from u to c_j and t'_o denotes the time that $\{u, c_j\}$ is traversed from c_j to u. Moreover,

$$q(u, c_j, t'_i, t'_o) = \begin{cases} w((u, c_j), t'_i) + w((u, c_j), t'_o) & \text{if } t'_i \neq t'_o \\ w((u, c_j), t'_i) & \text{otherwise} \end{cases}$$

If L is the total number of temporal edges, then f and g can be computed in time $O(nL^4)$ using dynamic programming. The optimal value of MTC is then $f(r, 0, \infty)$. In order to get

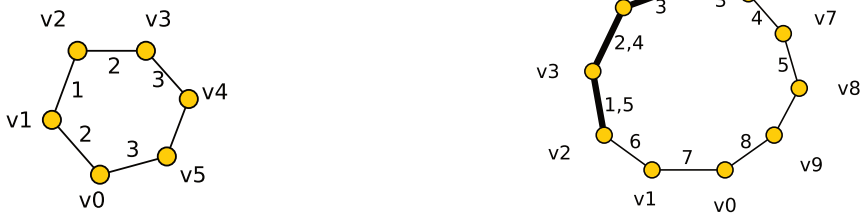

(a) v_1 is connected to any other vertex through temporal paths v_1, v_2, v_3, v_4 and v_1, v_0, v_5.

(b) The sector v_2, v_3, v_4, v_5. Its increasing and decreasing paths are $v_2, v_3, \ldots, v_8$ (using labels $1, 2, 3, 3, 4, 5$) and $v_5, v_4, \ldots, v_9$ (using labels $3, 4, 5, 6, 7, 8$) respectively.

Figure 2 Connectivity in temporal cycles.

the actual MTC solution, we just follow the parent links between states of the recurrence relation.

E.2 Temporal Connectivity on Cycles: 2-Approximation Algorithm

When the underlying graph is a cycle, we have a 2-approximation algorithm for MTC. In order for the temporal graph to be connected, for every vertex in the cycle there must exist temporal paths to every other vertex. Specifically, if we mark with $v_0, v_1, v_2, \ldots, v_{n-1}, v_n \equiv v_0$ the nodes of the cycle, then for a solution to be feasible, for every node v_i there must exist a node v_j, such that there are temporal paths $v_i, v_{i+1}, \ldots, v_{j-1}, v_j$ and $v_i, v_{i-1}, \ldots, v_{j+2}, v_{j+1}$ in the solution (vertex indices are taken modulo n). See also Figure 2.a.

▶ **Lemma E.2.** *Suppose that in some temporally connected subgraph $\mathcal{G}'$ of the temporal cycle, there is a temporal path $v_i, v_{i+1}, \ldots, v_j$ and for some $k \in \{i, i+1, \ldots, j\}$ there is also a temporal path $v_k, v_{k-1}, \ldots, v_{j+2}, v_{j+1}$. Then every vertex $v_p \in \{i, i+1, \ldots, k\}$ is connected to every other vertex in $\mathcal{G}'$.*

Proof. Since the temporal path $v_i, v_{i+1}, \ldots, v_j$ exists in $\mathcal{G}'$, so does the path $v_p, v_{p+1}, \ldots, v_j$. Also, since the path $v_k, v_{k-1}, \ldots, v_{j+2}, v_{j+1}$ exists in $\mathcal{G}'$, so does the path $v_p, v_{p-1}, \ldots, v_{j+2}, v_{j+1}$. So there are paths from v_p to every other vertex in $\mathcal{G}'$. ◀

This leads us to the following definition (see also Figure 2.b for an example).

▶ **Definition E.3.** Consider a temporal cycle $\mathcal{G}$ and a temporally connected subgraph $\mathcal{G}'$ of $\mathcal{G}$. A *sector* of $\mathcal{G}$ with respect to $\mathcal{G}'$ is a contiguous sequence of vertices $(v_i, v_{i+1}, \ldots, v_j)$ (it may be $i = j$) such that there exists a vertex $v_k \notin \{i, i+1, \ldots, j-1\}$ and the temporal paths $p_{\text{incr}} = (v_i, v_{i+1}, \ldots, v_k)$ and $p_{\text{decr}} = (v_j, v_{j-1}, \ldots, v_{k+1})$ are present in $\mathcal{G}'$. We refer to p_{incr} as the *increasing path* and to p_{decr} as the *decreasing path* of the sector $(v_i, v_{i+1}, \ldots, v_j)$. The *cost* of a sector $S = (v_i, v_{i+1}, \ldots, v_j)$ is the minimum, over all choices of increasing and decreasing paths pairs for S (in an increasing and decreasing path pair, if the increasing path ends at k, the decreasing ends at $k+1$), of the total cost of all temporal edges present in these increasing and decreasing paths.

Note that the vertices of a temporal cycle $\mathcal{G}$ with a temporally connected subgraph $\mathcal{G}'$ can always be partitioned into sectors with respect to $\mathcal{G}'$.

▶ **Lemma E.4.** *For a temporal cycle $\mathcal{G}$, there exists an optimal* MTC *solution $\mathcal{G}'$, such that no two sectors' increasing (resp. decreasing) paths use the same temporal edge (the same edge at the same time).*

Proof. Suppose that $\mathcal{G}'$ is partitioned to s disjoint sectors and that s is minimum among all optimal solutions. Furthermore, suppose that there are two (disjoint) sectors, $S_0 : v_0, v_1, \ldots, v_j$ and $S_1 : v_k, v_{k+1}, \ldots, v_l$, defined by this partition, whose increasing paths contain a common temporal edge. Note that disjointness implies $j < k$ and $l < n(\mathcal{G})$. We will show that this contradicts the minimality of s. Suppose the increasing path of sector S_0 is $v_0, v_1, \ldots, v_a$ and the increasing path of sector S_1 is $v_k, v_{k+1}, \ldots, v_b$. We distinguish between two cases:

$a \in [j, k-1]$: In order for the two increasing paths to have a common temporal edge, the increasing path of S_1 must be of the form $v_k, v_{k+1}, \ldots, v_z, v_{z+1}, \ldots, v_m$, where $z \in \{0, 1, \ldots, a-1\}$ and the edge between v_z and v_{z+1} is visited at the same time by the two increasing paths. This means that we can merge these two paths, creating a path $v_k, v_{k+1}, \ldots, v_a$. Combining this with the decreasing path of S_0, which is $v_j, v_{j-1}, \ldots, v_{a+1}$, this means that we can replace all sectors that contain vertices $v_k, v_{k+1} \ldots, v_j$ with a single sector consisting of exactly those vertices. Additionally, we have not added extra edges so the cost of the solution will not increase. Also note that the case $b \in [l, n(\mathcal{G}) - 1]$ is symmetric to the above case.

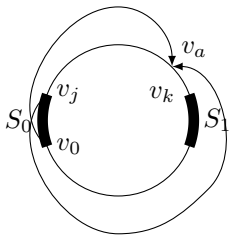

$a \in [k, n(\mathcal{G}) - 1]$ and $b \in [0, k-1]$: Let the common temporal edge of the two paths be connecting v_z and v_{z+1}. If $z \in \{0, 1, \ldots, k-1\}$, then merging the two increasing paths we obtain a new path $v_k, v_{k+1}, \ldots, v_z, v_{z+1}, \ldots, v_k, \ldots, v_a$, which means that we can replace all sectors with a single sector with no extra cost. As for the case $z \in \{k, k+1, \ldots, a-1\}$, we can combine temporal paths $v_0, v_1, \ldots, v_a$ with $v_l, v_{l-1}, \ldots, v_{a+1}$ to merge vertices $v_0, v_1, \ldots, v_l$ into a single sector with no extra cost.

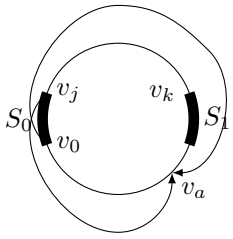

◀

The above lemma implies that in some MTC solution, every temporal edge is used at most twice by different sectors. Thus, ignoring the mutual dependence between sectors we can get a 2-approximate solution to our problem.

▶ **Theorem E.5.** *There is an algorithm that runs in polynomial time and outputs a feasible* MTC *solution of a temporal cycle $\mathcal{G}$, with cost at most 2 times the optimal.*

Proof. Denote the vertices of the cycle as $v_0, v_1, \ldots, v_n \equiv v_0$. First, we can try all possible vertices as the start of a sector. Let this be vertex v_0 of the cycle. Let $p(i)$ denote the partition of $v_0, v_1, \ldots, v_i$ into sectors with minimum cost (the cost of a set of disjoint sectors is defined as the sum of costs of the individual sectors) and $c(i)$ its cost.

Furthermore, define as inc_path(i, j) the minimum cost of a temporal path that visits $v_i, v_{i+1}, \ldots, v_j$ in this order and dec_path(i, j) the minimum cost of a temporal path that visits $v_i, v_{i-1}, \ldots, v_j$ in this order. inc_path and dec_path can be precomputed in time $O(nM)$ and space $O(n^2 + M)$, where M is the number of temporal edges.

Also, define sector_cost(i, j) to be the cost of the sector $v_i, v_{i+1}, \ldots, v_j$. This can be precomputed in time $O(n^3)$ and space $O(n^2)$, as

$$\text{sector_cost}(i,j) = \min_{k \in \{0,\ldots,n-1\} \setminus \{i, i+1, \ldots, j-1\}} \Big\{ \text{inc_path}(i,k) + \text{dec_path}(j, k+1) \Big\}$$

Then, we have that $c(0) = \text{sector_cost}(0, 0)$ and that for $i > 0$,

$$c(i) = \min \Big\{ \text{sector_cost}(0,i), \min_{j \in \{0,\ldots,i-1\}} \{c(j) + \text{sector_cost}(j+1, i)\} \Big\}$$

$p(i)$ can also be easily computed this way.

Finally, for the choice of v_0 that achieves the minimum $c(n-1)$, the algorithm outputs the union of the optimal increasing and decreasing paths of the sectors defined in $p(n-1)$.

The temporal graph that the algorithm produces is a feasible solution, because every vertex belongs to some sector and is thus connected to every other vertex.

To show that it is 2-approximate of the optimal solution, suppose that we double every temporal edge of the graph. The optimal cost OPT will not change. Yet, we can transform an optimal solution in which every temporal edge is used by at most 2 sectors to a solution in which every temporal edge is used by at most 1 sector in the doubled graph, with cost at most $2 \cdot OPT$. Either this or a smaller cost solution will be considered and produced by our algorithm. Overall, our algorithm uses $O(n^3 + n \cdot M)$ time and $O(n^2 + M)$ space. ◀